




Implementing an Open IBM SAN

Featuring the INRANGE Portfolio

Discover the latest additions to the IBM
SAN family

Enhance your skills while using
an easy-to-follow format

Grow with the new
technology



Jon Tate
Brian Cartwright
Sven Eichelbaum
Thomas Jahn

ibm.com/redbooks

Redbooks



International Technical Support Organization

**Implementing an Open IBM SAN
Featuring the INRANGE Portfolio**

December 2001

Take Note! Before using this information and the product it supports, be sure to read the general information in “Special notices” on page 369.

First Edition (December 2001)

This edition applies to those IBM SAN hardware and software products described herein.

Comments may be addressed to:
IBM Corporation, International Technical Support Organization
Dept. QXXE Building 80-E2
650 Harry Road
San Jose, California 95120-6099

When you send information to IBM, you grant IBM a non-exclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

© Copyright International Business Machines Corporation 2001. All rights reserved.

Note to U.S Government Users – Documentation related to restricted rights – Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract with IBM Corp.

Contents

Preface	ix
The team that wrote this redbook	ix
Special notice	xi
IBM trademarks	xii
Comments welcome	xii
 Chapter 1. Implementing Fibre Channel host adapter cards	1
1.1 Fibre Channel	2
1.1.1 Fibre Channel cables and adapter types	2
1.1.2 Fibre Channel node-to-node distances	3
1.1.3 LUN affinity	4
1.1.4 Targets and LUNs	4
1.2 Attaching an ESS to a FICON channel	4
1.2.1 Configuring the ESS for FICON attachment	4
1.2.2 Attachment considerations	5
1.2.3 Setting up ESCON and FICON links	5
1.2.4 Multipathing for ESCON and FICON	5
1.2.5 Attaching to a FICON channel or a FICON channel path group	6
1.2.6 Attaching an ESS to FICON channels	7
1.3 IBM iSeries FC host system attachment	8
1.3.1 Host limitations for the iSeries	8
1.3.2 Attachment requirements for the AS/400 or iSeries	9
1.3.3 Attachment considerations	9
1.3.4 Recommended configurations for the AS/400 or iSeries	10
1.3.5 Software requirements for the IBM OS/400 operating system	11
1.4 IBM RS/6000 or pSeries FC host system attachment	11
1.4.1 Attachment requirements	12
1.4.2 Installing the 2105 host attachment package	13
1.4.3 Replacing an older version of the 2105 installation package	13
1.4.4 Installing the 2105 host attachment package	14
1.4.5 Verifying the ESS configuration	15
1.4.6 ESS devices with multiple paths per LUN	15
1.5 IBM eServer xSeries or IBM NUMA-Q FC attachment	15
1.5.1 Attachment requirements	16
1.5.2 xSeries and NUMA-Q system requirements	16
1.5.3 Installing the IOC-0210-54 adapter card	17
1.5.4 Configuring the IOC-0210-54 adapter card	17
1.6 Compaq host system FC attachment	18

1.6.1	Attachment requirements	18
1.6.2	Compaq Tru64 UNIX Version 4.0x host system	19
1.6.3	Verifying the Compaq configuration	19
1.6.4	Operating system device recognition	19
1.6.5	Configuring AdvFS	22
1.6.6	Configuring devices to mount automatically	22
1.7	Hewlett Packard 9000 FC host system attachment	23
1.7.1	Attachment requirements	23
1.7.2	Installing the 2105 host install script file	24
1.7.3	Configuring the IBM 2105 for clustering	25
1.8	Novell NetWare FC host system attachment	26
1.8.1	Installing the QLogic QLA2100F adapter card	26
1.8.2	Installing the QLogic QLA2200F adapter card	27
1.8.3	Loading the current Fibre Channel adapter driver	29
1.8.4	Installing the Fibre Channel adapter drivers	29
1.8.5	ESS QLogic QLA2100F/2200F host adapter card	30
1.9	Sun FC host system attachment	30
1.9.1	Attachment requirements	30
1.9.2	Installing the Emulex LP8000 adapter card	32
1.9.3	Downloading the current Emulex Fibre Channel adapter driver	33
1.9.4	Installing the Fibre Channel adapter drivers	33
1.9.5	Installing the JNI PCI adapter card	35
1.9.6	Downloading the current JNI PCI Fibre Channel adapter driver	35
1.9.7	Installing the JNI PCI Fibre Channel adapter driver	36
1.9.8	Installing the JNI SBUS adapter card	36
1.9.9	Downloading current JNI SBUS Fibre Channel adapter driver	37
1.9.10	Installing the JNI SBUS Fibre Channel adapter driver	37
1.9.11	Installing the QLogic QLA2200F adapter card	38
1.9.12	Downloading the current Fibre Channel adapter driver	39
1.9.13	Installing the Fibre Channel adapter drivers	40
1.9.14	Configuring host device drivers	40
1.9.15	Tuning HBA configuration files	41
1.9.16	Installing the IBM Subsystem Device Driver	44
1.9.17	Setting the Sun host system parameters	45
1.10	Windows NT 4.0 FC host system attachment	47
1.10.1	Attachment requirements	47
1.10.2	Installing the QLogic QLA2100F adapter card	48
1.10.3	Installing the QLogic QLA2200F adapter card	49
1.10.4	Downloading the current Fibre Channel adapter driver	51
1.10.5	Installing the Fibre Channel adapter drivers	52
1.10.6	QLogic QLA2100F or QLogic QLA2200F configuration	52
1.10.7	Configuring for availability and recoverability	52
1.10.8	Setting the TimeOutValue registry	53

1.10.9	Installing Emulex LP8000 adapter cards.	53
1.10.10	Downloading the current Fibre Channel adapter driver	54
1.10.11	Installing the Fibre Channel adapter drivers	55
1.10.12	Parameter settings for the Emulex LP8000 on Windows NT	56
1.10.13	Verifying the configuration.	58
1.10.14	Verifying networking	58
1.11	Windows 2000 FC host system attachment	58
1.11.1	Attachment requirements	58
1.11.2	Installing the QLogic QLA2100F adapter card	59
1.11.3	Installing the QLogic QLA2200F adapter card	60
1.11.4	Downloading the current Fibre Channel adapter driver	62
1.11.5	Installing the Fibre Channel adapter drivers	63
1.11.6	QLogic QLA2100F or QLogic QLA2200F configuration	63
1.11.7	Configuring for availability and recoverability	64
1.11.8	Setting the TimeOutValue registry	64
1.11.9	Installing Emulex LP8000 adapter cards.	65
1.11.10	Downloading the current Fibre Channel adapter driver	65
1.11.11	Installing the Fibre Channel adapter drivers	66
1.11.12	Parameter settings for the Emulex LP8000 on Windows 2000. . .	67
1.11.13	Verifying the configuration.	69
1.11.14	Verifying networking	69
1.12	Locating the worldwide port name.	69
1.12.1	Fibre Channel port name identification	69
1.12.2	Locating the WWPN for an AS/400 iSeries.	70
1.12.3	Locating the WWPN for an IBM eServer or IBM NUMA-Q host . .	71
1.12.4	Locating the WWPN for an RS/6000 and pSeries host system . .	71
1.12.5	Locating the WWPN for a Compaq host system.	71
1.12.6	Locating the WWPN for a Hewlett Packard host system	72
1.12.7	Locating the WWPN for a Novell NetWare host system	73
1.12.8	Locating the WWPN for a Sun host system	73
1.12.9	Locating the WWPN for a Windows NT host system	74
1.12.10	Locating the WWPN for a Windows 2000 host system.	74
Chapter 2. IBM TotalStorage Enterprise Storage Server configuration . .		75
2.1	Introducing the ESS Web Interface.	77
2.1.1	Logging on to the ESS Specialist	78
2.1.2	Operating the ESS	82
2.1.3	Configuring the ESS	89
2.1.4	Administer the ESS.	93
2.2	Starting the Open Systems Storage Configuration	96
2.2.1	Viewing the storage configuration status	96
2.2.2	Using the Open System Storage panel.	104
2.3	Configuring disk groups.	107

2.3.1	Performing the disk group configuration	108
2.4	Creating and modifying host FC port definitions	116
2.4.1	Creating host FC port definitions.	117
2.4.2	Modifying a host FC port definition	123
2.5	Creating and assigning ESS logical volumes	126
2.5.1	Adding ESS logical volumes, step one	126
2.5.2	Adding ESS logical volumes — step two	129
2.6	Modifying ESS logical volume assignments	142
2.6.1	Sharing ESS logical volumes between two host FC ports	143
2.6.2	Removing volume assignments from host FC ports	151
2.7	Configuring ESS FC ports	158
2.7.1	Defining the topology for ESS FC ports	159
2.7.2	ESS managed ESS FC port WWPNS	164
2.7.3	Configuring host FC ports for FC-SW and FC-AL	167
 Chapter 3. Implementing the INRANGE FC/9000 Fibre Channel Director		175
3.1	Introduction to the INRANGE FC/9000 director	176
3.1.1	Product overview.	177
3.1.2	Currently supported zoning methodologies.	178
3.1.3	Management capabilities	178
3.1.4	Supported protocols	179
3.1.5	Supported host attachment	180
3.1.6	Supported device attachment	181
3.2	Installing an INRANGE SAN	181
3.2.1	Target SAN environment.	181
3.2.2	Verifying all prerequisites	182
3.2.3	Initial setup of INRANGE FC/9000 IP settings	183
3.2.4	Establishing network connection.	183
3.2.5	Installing the IN-VSN Enterprise Manager software	187
3.2.6	Accessing the management tool.	192
3.2.7	Defining users	197
3.2.8	Security considerations	205
3.3	Management of an INRANGE SAN.	207
3.3.1	Setting up operational parameters	207
3.3.2	Connecting to an INRANGE fabric	211
3.3.3	Assigning names and aliases	213
3.3.4	Setting the director clock.	218
3.3.5	Attaching loop ports	220
3.4	INRANGE zoning	237
3.4.1	Understanding INRANGE hard zoning	237
3.4.2	Understanding INRANGE broadcast zoning	241
3.4.3	Understanding INRANGE name server zoning.	242
3.4.4	Hard zones and name server zones together.	243

3.4.5	Defining hard zoning	246
3.4.6	Defining name server zones	252
3.5	Cascading of directors	263
3.6	Monitoring and maintenance of an INRANGE SAN	266
3.6.1	Management communication protocols	266
3.6.2	Microcode-loads	267
3.6.3	Monitoring user activities	267
3.6.4	Using the IN-VSN event log	269
3.7	ESS configuration guidelines	270
Chapter 4. Implementing the SAN Data Gateway		271
4.1	SAN Data Gateway	272
4.2	Installation	273
4.2.1	Setting the Ethernet address	275
4.2.2	Using Telnet on the SAN Data Gateway	275
4.2.3	Startup sequence	276
4.3	StorWatch SAN Data Gateway Specialist	277
4.3.1	Installing StorWatch Specialist	278
4.3.2	Using the StorWatch SAN Data Gateway Specialist	280
4.3.3	Upgrading the firmware	288
4.4	SCSI devices	294
4.4.1	LUN support	294
4.4.2	Device discovery	294
4.4.3	Assigning LUN IDs	295
4.4.4	Adding SCSI devices	295
4.5	SAN Data Gateway access options	296
4.5.1	Zoning	296
4.5.2	Virtual Private SAN	297
4.5.3	Combining Zoning and Virtual Private SAN	306
4.6	Adding Fibre Channel fabric components	307
4.6.1	Connecting an IBM SAN Fibre Channel Switch	307
4.7	High availability considerations	317
4.7.1	Single host with multiple Fibre Channel adapters	317
4.7.2	Multiple SCSI connections	318
4.7.3	Adding Fibre Channel switches	318
Chapter 5. Implementing the Vicom Fibre Channel SLIC Router		321
5.1	SLIC Router features	322
5.2	Installing the SLIC Router	323
5.3	SLIC Manager software	326
5.3.1	Installing the SLIC Manager software	327
5.3.2	Communicating to the Router	327
5.3.3	Starting the SLIC Manager	332

5.4 Using SLIC Manager	333
5.4.1 Drive properties	334
5.4.2 Router properties	334
5.4.3 Setting Router to master	335
5.4.4 The SignOn drive	336
5.5 Composite drive	337
5.5.1 Creating a composite drive	337
5.5.2 Composite drive properties	341
5.6 Mirror drive	342
5.6.1 Creating a mirror drive	343
5.6.2 Mirror drive properties	346
5.7 Instant Copy drive	347
5.7.1 Creating an Instant Copy drive	348
5.7.2 Instant copy drive properties	349
5.7.3 Adding an Instant Copy Drive to a mirror	350
5.7.4 Detach Instant Copy Drive from a mirror	352
5.8 Combining composite and mirroring	352
5.8.1 Creating a second composite drive	353
5.8.2 Creating the mirror	355
5.8.3 Viewing mirror drive using composite drives	356
5.9 Reusing logical drives	357
5.9.1 Remove a logical drive	357
5.9.2 Mapping a general spare	359
5.9.3 Removing a mirror containing composite drive	359
5.10 Expanding the SLIC system	361
5.10.1 Adding disk	361
5.10.2 Adding Routers	361
5.10.3 Adding hosts	364
Related publications	365
IBM Redbooks	365
Other resources	365
Referenced Web sites	367
How to get IBM Redbooks	368
IBM Redbooks collections	368
Special notices	369
Glossary	371
Index	383

Preface

“Do everything that is necessary and absolutely nothing that is not.”

In this IBM Redbook, we have tried to consolidate as much of the critical information as possible while covering procedures and tasks that are likely to be encountered on a daily basis.

Each of the products described has much, much more functionality than we could ever hope to cover in just one redbook. The IBM SAN portfolio is rich in quality products that bring a vast amount of technicality and vitality to the SAN world. Their inclusion and selection is based on a thorough understanding of the storage networking environment that positions IBM, and therefore its customers and partners, in an ideal position to take advantage by their deployment.

In this redbook we cover the latest additions to the IBM SAN family, which includes products from companies such as Brocade, INRANGE, and McDATA. We show how they can be implemented in an open systems environment, and we focus on the Fibre Channel protocol (FCP) environment in particular. We address some of the key concepts that they bring to the market, and in each case, we give an overview of those functions that are essential to building a robust SAN environment.

The team that wrote this redbook

This redbook was produced by a team of specialists from around the world working at the International Technical Support Organization, San Jose Center.

Jon Tate is a Project Leader for SAN TotalStorage Solutions with the International Technical Support Organization, San Jose Center. Before joining the ITSO in 1999, he worked in the IBM Technical Support Center, providing level 2 support for IBM storage products. Jon has 16 years of experience in storage software and management, services, and support, and he is an IBM SAN Certified Specialist.

Brian Cartwright is an Advisory I/T Specialist with the Storage Systems Group and is based in Brisbane, Australia. Brian has 14 years of experience in the I/T industry with 6 years at IBM working as a Technical Specialist in the S/390 and Storage Systems Groups. During this time, Brian has been providing pre-sales and technical support for IBM Storage Solutions and has been heavily involved in

the design and implementation of IBM Storage Area Networks across a wide customer base. Brian's expertise includes SAN design, implementation, and management, and covers a wide range of disk and tape solutions across a number of different operating system platforms.

Sven Eichelbaum is an I/T Specialist at StorAdvantage GmbH, an IBM Business Partner based in Germany which specializes in SAN solutions. He has 4 years of experience in designing and implementing storage architectures. Before joining StorAdvantage in 2001 he worked in IBM SSD Germany as System Engineer. He holds a Dipl. Betriebswirt degree in Computer Sciences for Economy from the Staatliche Studienakademie Sachsen.

Thomas Jahn is an I/T Specialist with IBM Germany SSG. He has 4 years of experience providing technical support in IBM. Thomas has provided technical support for networking and server consolidation on OS/390 UNIX for IBM and its customers before he joined SSD two years ago. He is currently engaged in providing technical support for open systems storage solutions across multiple platforms and customer bases. He holds a Dipl. Ing. degree in Computer Science from the Staatliche Studienakademie Sachsen.

Thanks to the following people for their invaluable contributions to this project:

Scott Drummond
IBM Storage Subsystems Division

Emma Jacobs
Yvonne Lyon
Deanna Polm
Sokkieng Wang
International Technical Support Organization, San Jose Center

Henry Caudillo
Mac Diller
Glenda Fuller
Mike Hrencecin
Ronda Hruby
Robert Moon
Tammy Sokol
Peter Thurston
Diana Tseng
Karen Ward
Rainer Wolafka
Michelle Wright
John Young
Ruoyi Zhou
IBM Storage Subsystems Division

Richio Aikawa
Jon Krueger
Emulex Corporation

Dave Burchwell
Jack Consoli
Richard Kurzban
Mike Naylor
INRANGE Technologies Corporation

Chris Burke
JNI Corporation

Rob Jones
QLogic Corporation
Philipp Alexander
StorAdvantage


Special notice


This publication is intended to help systems and storage administrators install IBM SAN portfolio equipment. The information in this publication is not intended as the specification of any programming interfaces that are provided by any of the SAN hardware and software components contained herein. See the PUBLICATIONS section of the IBM Programming Announcement for the SAN hardware and software components contained herein for more information about what publications are considered to be product documentation.

Attention: What this means specifically for this redbook is that as we have been working with Beta code, some of the screen captures may be different from products that appear on the market today.

IBM trademarks

The following terms are trademarks of the International Business Machines Corporation in the United States and/or other countries:

e (logo)® 
IBM ®

Redbooks
Redbooks Logo 

AIX
AS/400
BookMaster
Current
Early
Enterprise Storage Server
ESCON
FICON
FlashCopy
iSeries
Magstar
Netfinity
Notes
NUMA-Q

OS/390
OS/400
pSeries
PTX
RS/6000
S/390
SANergy
SP
StorWatch
System/390
TotalStorage
Wave
xSeries
zSeries

Comments welcome

Your comments are important to us!

We want our IBM Redbooks to be as helpful as possible. Send us your comments about this or other Redbooks in one of the following ways:

- ▶ Use the online **Contact us** review redbook form found at:
ibm.com/redbooks
- ▶ Send your comments in an Internet note to:
redbook@us.ibm.com
- ▶ Mail your comments to the address on page ii.



Implementing Fibre Channel host adapter cards

In this chapter we describe the steps involved in implementing Fibre Channel host adapter cards in an IBM TotalStorage Enterprise Storage Server (ESS) environment.

Fibre Channel is a 100-MBps, full-duplex, serial communications technology to interconnect I/O devices and host systems that are separated by tens of kilometers.

1.1 Fibre Channel

Fibre Channel transfers information between the sources and the users of the information. This information can include commands, controls, files, graphics, video, and sound. Fibre Channel connections are established between Fibre Channel ports that reside in I/O devices, host systems, and the network that interconnect them. The network consists of elements like switches, hubs, bridges, and repeaters that are used to interconnect the Fibre Channel ports. The ESS architecture supports three basic topologies:

- ▶ Point-to-point
- ▶ Switched fabric
- ▶ Arbitrated loop

Before you start, check that the Fibre Channel host has the correct Fibre Channel host bus adapters installed or at least has slots available for them. For a list of operating systems and the host bus adapters for Fibre Channel attachment, see following the Web site:

www.storage.ibm.com/hardsoft/products/ess/supserver.htm

1.1.1 Fibre Channel cables and adapter types

Use feature code 3022 for the short-wave, Fibre Channel host adapter card. Feature code 3022 comes with a complimentary 31 m (101 ft), 50 micron multimode fibre cable.

You can also use a 62.5-micron fiber-optic cable with the short-wave Fibre Channel card. For a 50-micron multimode fiber-optic cable, use a cable length of 2 m - 500 m (6.5 ft - 1640 ft).

For a 62.5-micron multimode fiber-optic cable, use a cable length of 2 m - 175 m (6.5 ft - 575 ft).

If you use Fibre Channel to ESCON cable converters at both ends of the link to allow connection through an ESCON cable, the ESCON cable must be a 62.5-micron fiber-optic cable. The maximum length for the 62.5-micron fiber-optic cable is 175 m (575 feet).

Limitations on distances between nodes on a Fibre Channel network might exist, depending on the upper-layer protocol (ULP).

1.1.2 Fibre Channel node-to-node distances

You can order and have IBM install a maximum of 16 Fibre Channel adapters on ESS Models F10 and F20. Each adapter has a single host port. You can install both long-wave adapters and short-wave adapters in a single ESS.

Table 1-1 lists the cables that are available for the long-wave and short-wave adapters. This table also lists the distances that the cables support. See your IBM field office about 62.5 micron cables or longer cables.

- Feature code 3021 — Fibre Channel long-wave host adapter (optional 1-16) for open-systems, AS/400, iSeries, S/390, and zSeries hosts. This adapter includes a 31 m (101 ft), long-wave optics (9 micron) cable, P/N 08H2774 with an SC-type connector. It provides an interface that supports 100 Mbps full-duplex data transfer.
- Feature code 3023 — Fibre Channel short-wave host adapter (optional 1-16) for open-systems, AS/400, iSeries, S/390, and zSeries hosts. This adapter includes 31 m (101 ft), 50 micron cable, P/N 54G3384 with an SC-type connector. It provides an interface that supports 100 Mbps full-duplex data transfer.

Table 1-1 *Fibre Channel cables*

Adapter	Cable Type	Distance
FC 3023 (short wave)	50 micron, multimode (500 MHz/km)	500 m (1635 ft)
	62.5 micron, multimode (200 MHz/km)	300 m (984 ft)
	62.5 micron, multimode (160 MHz/km)	250 m (817 ft)
FC 3021 (long wave)	9 micron, singlemode	50 km (31 mi)
	62.5 micron, multimode (500 MHz/km)	550 m (1799 ft)
	62.5 micron, multimode (200 MHz/km)	550 m (1799 ft)
	62.5 micron, multimode (160 MHz/km)	550 m (1799 ft)

For Fibre Channel, the maximum distance between fabric switches, fabric hubs, and link extenders and a host Fibre Channel port or an ESS Fibre Channel port is limited to 11 km (6 mi). The maximum distance might be greater than 11 km (6 mi) when a link extender provides appropriate target initiator or controller emulation functions such that the length of the connection perceived by the ESS does not exceed 11 km (6 mi).

Link extenders with emulation functions should not be used on links over which synchronous PPRC operations are performed because of the additional path delay introduced by these units.

1.1.3 LUN affinity

For Fibre Channel attachment, LUNs have an affinity to the host's Fibre Channel adapter through the world-wide port name (WWPN) for the host adapter. In a switched fabric configuration, a single Fibre Channel host could have physical access to multiple Fibre Channel ports on the ESS. In this case, you can configure the ESS to allow the host to use either:

- ▶ All physically accessible Fibre Channel ports on the ESS
- ▶ Only a subset of the physically accessible Fibre Channel ports on the ESS

In either case, the set of LUNs that are accessed by the Fibre Channel host are the same on each of the ESS ports that can be used by that host.

1.1.4 Targets and LUNs

For Fibre Channel attachment, each Fibre Channel host adapter can architecturally attach up to 256 LUNs. The ESS supports only a maximum of 4096 LUNs divided into a maximum of 16 logical subsystems each with up to 256 LUNs. If the software in the Fibre Channel host supports the SCSI command Report LUNs, then you can configure all 4096 LUNs on the ESS to be accessible by that host. Otherwise, you can configure no more than 256 of the LUNs in the ESS to be accessible by that host.

1.2 Attaching an ESS to a FICON channel

This topic tells you how to configure the IBM Enterprise Storage Server for FICON attachment.

1.2.1 Configuring the ESS for FICON attachment

You can perform a FICON channel attachment on the IBM 2105 Model F10 and F20. You cannot perform a FICON channel attachment on the IBM 2105 Model E10 and E20. When you attach an F10 or F20 to a FICON interface, you must use the following host adapter feature codes:

- ▶ Feature code 3021 — This is a long wave laser adapter that includes a 31-m (100-ft) 50-micron cable with duplex connectors. Because the 3021 uses one of four slots in one of the four I/O bays, you can have a maximum of 16 adapters in the ESS. This allows you to have a maximum of 16 FICON interface attachments.

- Feature code 3023 — This is a short wave laser adapter that includes a 31-m (100-ft) (50-micron) cable with duplex connectors. When you use the 3023 adapter to perform a FICON channel attachment, you must use the cables for FICON. The adapters cannot be shared with any Fibre Channel protocol attachments.

Because the 3023 uses one of four slots in one of the four I/O bays, you can have a maximum of 16 adapters in the ESS. This allows you to have a maximum of 16 FICON interface attachments. If the attachments are all point-to-point, you can attach directly to 16 FICON channels. If you attach to a switch or director, you can attach a maximum of 128 FICON channels per ESS FICON adapter. For this scenario, you must be able to attach to a number of hosts. The ESS allows 256 logical paths per FICON link (compared to just 64 for ESCON), 128 logical paths per logical subsystem, and 2048 logical paths for each ESS.

Before FICON, you could only do a Fibre Channel connect for Fibre Channel protocol with feature code 3022. Feature code 3023, with a 9-micron single mode fiber-optic cable, increases the point-to-point distance from 500 m to 10 km.

The increased distance provides greater configuration options with IBM S/390 or zSeries processors with FICON host adapters.

1.2.2 Attachment considerations

This topic describes some things you should consider before you configure your system with a FICON interface.

1.2.3 Setting up ESCON and FICON links

If the system requires x ESCON links, where x is the number of links to get the performance and availability attributes you want, you must consider the number of ESCON you need. For example, you can map four ESCON links to a single FICON link and maintain approximately equivalent performance. If the ESCON channel use is low, you can map six or eight ESCON links to a single FICON link.

1.2.4 Multipathing for ESCON and FICON

Consider the difference between the path groups when you compare FICON to ESCON. For example, for ESCON, you can configure four or eight paths per path group from a host to an ESS. For ESCON, you want at least four paths in the path group to maximize performance. Most ESCON controllers implement channel command execution that partially synchronizes the lower DASD interface with the upper channel interface. This channel command only allows you a very short time to reconnect. The consequence is reconnections that fail.

When you have eight paths in the path group, it minimizes the number of missed reconnections. Increasing the number of path groups does not minimize the number of missed reconnections substantially. If you use eight paths in path groups, you can increase the overall throughput.

For FICON controllers, there is no synchronization between the lower DASD interface and the upper channel interface. The number of paths in the path group depend on the throughput requirement. If it takes x paths to satisfy the throughput requirement, where x is the number of paths, set the path group to x .

Note: x must be a minimum of two and cannot exceed a maximum of eight.

1.2.5 Attaching to a FICON channel or a FICON channel path group

When you attach multiple controllers to a channel, you can use a switch (director) for each controller, or an ESCON or FICON channel that has a direct connection to the channel. I/O does not flow through all the other controllers before you get to the target controller. I/O goes directly to the target controller. When multiple controllers are connected to a channel through a switch, you create the logical equivalent of the parallel interconnection.

With the parallel interface and with the ESCON interface, the channel and controller communicate to form a private connection. None of the other controllers on the channel can communicate with the channel while this private connection is in place. The private connection supports input and output between the channel and the controller. It can run slowly, depending upon the factors that affect the controller and the device. The protocol does not allow any of the serially connected controllers to use any spare cycles. The result is poor performance.

FICON does not support a private connection. FICON performs frame (or packet) multiplexing. A configuration with the serially connected controllers communicates with the controllers simultaneously. It can multiplex I/O operations across all controllers simultaneously. No interface cycles are wasted because of the private connection. You can serially connect controllers with FICON.

The next question though is whether or not it is OK to serially connect DASD control units with tape controllers. Tape generally performs much larger I/O operations at any instant in time. Therefore, even with FICON, when you have Tape I/O running, you can temporarily lockout some DASD I/O. Hence, it is still better not to put tape and DASD on the same FICON channel.

1.2.6 Attaching an ESS to FICON channels

You can use the following FICON adapters with the IBM S/390 Enterprise Storage Servers.

- ▶ Feature code 2314 — This is the long-wave laser adapter.
- ▶ Feature code 2316 — This is the short-wave laser adapter.

You can use the following FICON adapters with the IBM S/390 or zSeries system:

- ▶ Feature code 2315 — This is the FICON long-wave laser adapter. This adapter has two ports per adapter. This adapter is a 9-micron single mode cable, but you can use it with a 62.5-micron multimode cable when you attach mode-conditioning cables at each end.
- ▶ Feature code 2318 — This is the FICON short-wave laser adapter. This adapter has two ports per adapter. The short-wave laser adapter supports the 50 and 62.5-micron multimode cable.

You can attach the FICON channels directly to an ESS or you can attach the FICON channels to a Fibre Channel switch. When you attach the FICON channels directly to an ESS, the maximum number of FICON attachments is 16. That is the maximum number of host adapters you can configure in an ESS. When you use an ESS host adapter to attach to FICON channels either directly or through a switch, the adapter is dedicated to FICON attachment and may not be simultaneously attached to Fibre Channel protocol hosts.

When you attach an ESS to FICON channels through one or more switches, the maximum number of FICON attachments is 128 per ESS adapter. The directors provide very high availability with redundant components and no single points of failure or repair.

You can use the INRANGE FC/9000 Fibre Channel director to attach Fibre Channel protocol hosts and devices in addition to the FICON hosts and devices. For these configurations, the Fibre Channel protocol hosts should communicate only with the Fibre Channel protocol devices. The FICON hosts should communicate only with the FICON devices. IBM recommends that you set up zones in the directors to guarantee that none of the Fibre Channel protocol hosts or devices can affect the FICON traffic.

When you attach FICON products to switches or directors, note that switch cascading is not allowed. You cannot configure a fabric of multiple interconnected directors and have a FICON channel attached to one director communicate to a FICON control unit attached to another director.

The FICON architecture prohibits this capability. The reason for the restriction is because the base S/390 and zSeries I/O architecture uses a single byte for addressing the I/O devices. This one-byte I/O address is not compatible with the Fibre Channel, 3-byte port address. The FICON solution to this problem is to disallow switch cascading.

1.3 IBM iSeries FC host system attachment

This topic describes the host system requirements and provides the procedure to attach your IBM AS/400 or iSeries host system to the IBM Enterprise Storage Server (ESS) with Fibre Channel adapters.

Note: The IBM AS/400 or iSeries host systems do not support serially interconnected host systems attached to the ESS. They also do not support serially interconnected ESS attachments to a single port on the host adapter.

Your IBM AS/400 or iSeries host systems support the ESS as a peripheral device.

The ESS presents individual 2105 LUNs to the AS/400 or iSeries host system. You do not need to perform manual tasks to assign LUN addresses, as the ESS will do it automatically during configuration.

Note: You cannot use the IBM Subsystem Device Driver on the AS/400 or iSeries host system.

1.3.1 Host limitations for the iSeries

In Table 1-2 we show a description of the LUN assignments for the iSeries host system.

Table 1-2 iSeries host system limitations

Host system	LUN assignments per target	Configuration notes
iSeries (Fibre Channel)	0 - 32	There is one target per iSeries adapter

The naming convention for the iSeries now describes models 270 and 8xx, which attach through an FC adapter 2766. Use migration tower 5035 or 5077 for the model 8xx. You can also make a connection using a 6501 with SCSI cables.

1.3.2 Attachment requirements for the AS/400 or iSeries

This section lists the requirements for attaching the ESS to your host system:

1. Ensure that you have all of the items in the equipment list.
2. Obtain the documents for the IBM AS/400 or iSeries host system from the following Web site:
`publib.boulder.ibm.com/pubs/html/as400/infocenter.htm`
3. See the following Web site for details about program temporary fixes (PTFs) that you need to install on your AS/400 or iSeries host system:
`www.storage.ibm.com/hardsoft/products/ess/supserver.htm`
4. Check the LUN limitations for your host system.
5. Contact your IBM service support representative to install and configure the IBM ESS.

1.3.3 Attachment considerations

This section lists the attachment considerations for an AS/400 or iSeries host system. The following must be taken into account:

- ▶ You can specify a LUN serial number that is eight characters. For example, you can specify a LUN serial number, 0LLLLNNN, where:

The LUN serial number is eight characters in the format 0LLLLNNN, where:

LLLL — a unique volume number assigned by the ESS when the LUN is created.

NNN — low-order three characters of the ESS unit serial number or unique three character value entered using an ESS service panel menu option.

The menu option is found by using the following method:

- a. From Service Menu choose the option Configuration Options Menu
- b. Take the option Change / Show Control Switches.
- c. Highlight the option AS/400 LUN Serial Number Suffix.
- d. On the option Control Switch Value press F4 to list the available values.

Note: The default value is the last three digits of the serial number for the ESS. This might create problems on iSeries systems that have multiple ESSs attached to the same host system. We recommended that you use the value **Use_Box_S/N**.

The following considerations need to be taken into account:

- ▶ You cannot specify a LUN size of 4.190-GB LUN for the SCSI Fibre Channel protocol (FCP) attachment.
- ▶ You can specify 1- 32 LUNs for each attachment to an AS/400 or iSeries Fibre Channel protocol adapter.
- ▶ Fibre Channel attached LUNS are identified as the 2105 device type on the iSeries host system.
- ▶ You can place an AS/400 or iSeries volume in the ESS storage arrays according to the selected host system attachment type.

For a Fibre Channel attachment, you must place the volumes in ESS RAID-5 storage arrays that have capacity available. You can spread the volumes across arrays that are attached to multiple device adapters.

- ▶ You cannot place an AS/400 or iSeries volumes in an ESS non-RAID storage array.
- ▶ You cannot share an AS/400 or iSeries volume with more than one Fibre Channel system attachment.
- ▶ The number of volumes that you can create is determined by the attachment type and by the available ESS storage array capacity.
- ▶ You can create 1 - 32 LUNs for a Fibre Channel attachment.

1.3.4 Recommended configurations for the AS/400 or iSeries

IBM recommends the following configurations:

- ▶ Install the 2766 adapter card in the AS/400 iSeries system unit or the HSL PCI I/O towers.

Feature code 2766 is an IOA (I/O adapter).

Important: Do not use more than two 2766 adapter cards for each AS/400 or iSeries tower.

The following must be taken into account:

1. Only one 2766 adapter is supported per I/O processor (IOP) and requires a dedicated IOP. No other I/O adapters are supported under the same IOP.
2. Only two 2766 adapters are supported per a multi-adapter bridge.

1.3.5 Software requirements for the IBM OS/400 operating system

Release V3R1.0 and later releases of the OS/400 or iSeries operating system support 4-GB and 8-GB capacity disk drives. In Table 1-3 we list the required program temporary fixes (PTFs) for OS/400 or iSeries versions that do not include the PTFs in the base code.

Table 1-3 OS/400 program temporary fixes for V3R1.0 through V4R1.4

Operating System Version Level	PTF
V3R1.0	SF44131
V3R2.0	SF44132
V3R6.0	SF44126
V3R7.0	SF44127
V4R1.0	SF44113
V4R1.4	SF44745

Releases V4R2.0, V4R3.0, and V4R4.0 also support 18-GB and 36-GB capacity logical unit numbers with the PTFs shown in Table 1-4.

Table 1-4 OS/400 program temporary fixes for V4R2.0 and higher

Operating System Version Level	PTF
V4R2.0	SF44114 MF21961 MF21975
V4R3.0	MF21962 MF21979
V4R4.0	MF21963 MF21978

1.4 IBM RS/6000 or pSeries FC host system attachment

This topic describes the host system requirements and provides the procedures to attach an ESS to any of the following:

- ▶ IBM RS/6000
- ▶ pSeries host system
- ▶ IBM RS/6000 Series Parallel (SP) Complex
- ▶ pSeries SP Complex host system

Note: For an RS/6000 and pSeries host system, you can use two topologies:

- ▶ Point-to-point (switched fabric) topology
- ▶ Arbitrated loop topology

The RS/6000 and pSeries host system does not support more than one host bus adapter on the loop. The RS/6000 and pSeries host system does support a direct connection of the RS/6000 and pSeries host system to an ESS using the Fibre Channel arbitrated loop protocol.

1.4.1 Attachment requirements

This section lists the requirements for attaching the ESS to your host system.

- ▶ Ensure that you have all of the items listed in the Equipment requirements.
- ▶ Ensure that you have the installation script files. These files are on the diskette and the compact disc you receive with the ESS.
- ▶ Ensure that you have 1 MB minimum of hard disk space available to install the AIX host attachment package.
- ▶ Ensure that you have the documentation for your host system and the IBM Enterprise Storage Server User's Guide. The User's Guide is on the compact disc that you receive with the ESS.

For details about the release level for your operating system, see the following Web site:

www.storage.ibm.com/hardsoft/products/ess/supserver.htm

Either you or an IBM service support representative (SSR) must perform the following tasks to install and configure an ESS.

1. The IBM SSR installs the ESS by using the procedures in the IBM Enterprise Storage Server Service Guide.
2. Either you or an IBM SSR defines the Fiber Channel port configuration if you did not do it during the installation of the ESS or Fibre Channel adapters.

Note: Use the information on the logical configuration work sheet in the IBM Enterprise Storage Server Configuration Planner that you previously filled out.

3. Either you or an IBM SSR configures the host system for the ESS. Use the instructions in your host system publications.

4. Either you or an IBM SSR checks the LUN limitations for the RS/6000 and pSeries. See Table 2.

Note: The IBM Subsystem Device Driver supports RS/6000 and pSeries host systems in a clustering environment. To have failover protection on an open system, the IBM Subsystem Device Driver requires a minimum of 2 Fibre Channel adapters. The maximum number of Fibre Channel adapters supported is 16 for a total of 16 Fibre Channel ports.

1.4.2 Installing the 2105 host attachment package

This section provides the instructions to install the host attachment package for the ESS on your host system. IBM recommends that you run the host attachment package on each host system attached to the ESS for which an installation script is provided.

Before installing the 2105 host attachment package

Perform the following steps before you install the host attachment package:

1. Attach the ESS to your host system. See Attaching an ESS to an open-systems host with SCSI adapters.
2. Turn on the host system and all attachments.
3. Ensure that you have root access.
4. Ensure that you have administrator knowledge.
5. Ensure that you have knowledge of the System Management Interface Tool (SMIT).

1.4.3 Replacing an older version of the 2105 installation package

If you want to replace an older version of the host attachment package (tar version) and have data that exists on all configured 2105 disks, the code prompts you to remove all ESS product-related hdisk devices. Perform the following steps to remove the devices:

1. Run the **umount** command on the file system.
For example, type `umount -t x`, where `x` is the file system name.
2. Run the **varyoffvg** 2105 command for the volume group.
For example, type `varyoffvg -s VGname`
3. Type **rmdev -d1** on the command line to unconfigure the 2105 devices.

After you install the `ibm2105.rte` file and all of the 2105 devices are reconfigured, vary on the volume groups and remount the file systems. The data on the file systems should be available again.

1.4.4 Installing the 2105 host attachment package

Perform the following steps by using SMIT to install the IBM 2105 host attachment on your system. Install the host attachment package from a compact disc or a diskette. You must have superuser authority to complete the instructions.

Note: The following procedure is an example. The example uses `/dev/cd0` for the address of the CD-ROM drive. Your address might be different.

1. From your desktop window, type `smit install_update` to go directly to the installation panel.
2. Click **Install** and **Update** from the Latest Available Software and press Enter.
3. Press F4 to open the Input Device/Directory for Software window.
4. Select the CD-ROM drive that you are using for the installation, for example, `/dev/cd0`.
5. Press Enter.

The Install and Update from the Latest Available Software window opens.

6. Click **Software** to Install and press F4.
7. Select **Software Packages** and press F7.

The Install and Update from the Latest Available Software panel is displayed with the name of the software you selected to install.

8. Check the default option settings to ensure that they are what you need.
9. Press Enter to install the software.

SMIT responds with the following question: Are you sure?

10. Press Enter to continue.

The installation process may take several minutes. A message is displayed when the installation process is complete.

11. Press Enter to continue. The installation process may take several minutes.
A message is displayed when the installation process is complete.
12. Press F10 when the installation process is complete.
13. Exit from SMIT.

14. Remove the compact disc.
15. Shut down the host system.
16. Turn on the host system.

1.4.5 Verifying the ESS configuration

To verify the configuration of the ESS on the AIX host system, type the following command:

```
lsdev -Cc disk | grep 2105
```

A list of all IBM ESS devices is displayed as shown in the following example:

```
hdisk3 Available 30-68-01 IBM FC2105F20
hdisk4 Available 30-68-01 IBM FC2105F20
hdisk5 Available 30-68-01 IBM FC2105F20
```

These indicate that the installation is successful.

If a device is listed as another type of device, the message shown below is displayed. This message indicates that the configuration was not successful.

```
hdisk3 Available 30-68-01, Other FCSCSI disk device
hdisk4 Available 30-68-01, Other FCSCSI disk device
hdisk5 Available 30-68-01, Other FCSCSI disk device
```

1.4.6 ESS devices with multiple paths per LUN

The ESS supports multiple path configurations for a LUN. This means that you can have multiple hdisks available on the AIX server for each physical LUN. If you create a PVID sector 0 of a LUN and you delete all hdisks from the system with the **rmdev** command, you must restart the system. If you want to restore all multiple paths for all LUNS, use the **cfgmgr** command for each Fibre Channel adapter.

1.5 IBM eServer xSeries or IBM NUMA-Q FC attachment

This topic tells you how to attach an ESS to an xSeries or an IBM NUMA-Q host system with Fibre Channel adapters. We also tell you how to install and configure the IOC-0210-54 adapter card.

Note: You must use the switched fabric topology to attach the ESS to either an xSeries or to an IBM NUMA-Q host system.

The ESS also offers the following interim support for Fibre Channel attachment:

Feature Code 3019 — With the feature code 3019, you can attach an ESS to an xSeries or an IBM NUMA-Q host system through the NUMA-Q Fibre Channel- to-SCSI bridge. This feature code includes one SCSI adapter that you purchase and a no-cost loan of a NUMA-Q Fibre Channel to SCSI bridge. IBM requires that you sign a loan agreement for the bridge.

Note: Feature code 3019 is not a standard feature. To get feature code 3019, contact your IBM sales representative.

For more information about how to attach an xSeries or a NUMA-Q host system with Fibre Channel adapters, see the NUMA-Q ESS Integration Release Notes and the Fibre Channel Subsystem Installation Guide. To obtain a copy, see your IBM sales representative.

1.5.1 Attachment requirements

This section lists the requirements for attaching the ESS to your host system.

Either you or an IBM service support representative (SSR) must perform the following tasks to install and configure an ESS:

1. The IBM SSR installs the IBM ESS by using the procedures in the IBM Enterprise Storage Server Service Guide.
2. Either you or the IBM SSR defines the Fibre Channel host system with the worldwide port name identifiers. For information about how to locate the worldwide port name for an xSeries or an NUMA-Q host system, see 1.12, “Locating the worldwide port name” on page 69.
3. Either you or the IBM SSR defines the Fiber Channel port configuration if you did not do it during the installation of the ESS or Fibre Channel adapters.

Note: Use the information on the logical configuration work sheet in the IBM Enterprise Storage Server Configuration Planner that you previously filled out.

4. Configure your host system for the ESS by using the instructions in your host system publications.

1.5.2 xSeries and NUMA-Q system requirements

The ESS is supported on the xSeries and the NUMA-Q host systems by a module of code that is incorporated into Service Pack 3 for PTX V4.5.2. To install Service Pack 3:

1. Insert the Service Pack 3 compact disc into the CD-ROM or CD-R drive.
2. Open the README file for instructions on installing Service Pack 3.

In Table 1-5 we show the NUMA-Q system requirements. Support for Copy Services on PTX V4.5.2 requires a special Technology Pack. You can obtain the Technology Pack through an IBM sales representative who handles your xSeries and NUMA-Q purchases.

Table 1-5 xSeries and NUMA-Q system requirements for the ESS

xSeries and NUMA-Q System Element	Requirement
PTX operating system	Version 4.5.2 or higher
NUMA-Q and xSeries hardware models	All NUMA-Q and xSeries Quad-based systems
Fibre Channel host adapter	Emulex LP7000E with firmware SF 3.2.1
Fibre Channel switch	IBM 2109 Model S08 or IBM 2109 Model S16
Clustered NUMA-Q and xSeries hosts	ptx/Clusters V2.2.1

1.5.3 Installing the IOC-0210-54 adapter card

Perform the following steps to install the IOC-0210-54 adapter card:

1. Contact your IBM SSR to install the IOC-0210-54 adapter card in the ESS.
2. Connect the cable to the ESS port.

The SSR establishes the private LAN connection between both clusters on the ESS, the Ethernet hub, and the ESS personal computer console.

Preconfigured multimode optical cables are available to connect the ESS to the NUMA-Q host system. The 8m (24-ft) cable might be required. The way you connect the cable to the ESS through the Fibre Channel switch depends on the level of I/O throughput.

For information about connection schemes, see the Fibre Channel Subsystems Installation Guide at the following Web site:

techdocs.sequent.com/staticpath/shsvccd/start.htm

3. Restart the server.

1.5.4 Configuring the IOC-0210-54 adapter card

To configure the IOC-0210-54 adapter card, contact your IBM SSR or see the IBM Enterprise Storage Server Web Interface User's Guide.

1.6 Compaq host system FC attachment

This topic describes the host system requirements and provides the procedure to attach a Compaq Alpha server to an ESS. The ESS supports the following Compaq Alpha server models:

- ▶ 2100
- ▶ 4100
- ▶ 8200
- ▶ 8400

1.6.1 Attachment requirements

This section lists the requirements to attach the ESS to your host system.

- ▶ Ensure that you have all of the items listed in Equipment requirements.
- ▶ Check the logical unit number limitations for your host system. See Table 2.
- ▶ Ensure that you have the documentation for your host system and the IBM Enterprise Storage Server User's Guide. The User's Guide is on the compact disc that you receive with the ESS.

See the following Web site for details about the release level for your operating system:

www.storage.ibm.com/hardsoft/products/ess/supserver.htm

Either you or an IBM service support representative (SSR) must perform the following tasks to install and configure an IBM ESS.

1. An IBM SSR installs the IBM Enterprise Storage Server by using the procedures in the IBM Enterprise Storage Server Service Guide.
2. Either you or an IBM SSR assigns the Fibre Channel hosts to the Fibre Channel ports on the ESS.

Note: Use the information on the logical configuration work sheet in the IBM Enterprise Storage Server Configuration Planner that you should have previously filled out.

3. Either you or an IBM SSR configures the host system for the ESS. Use the instructions in your host system publications.

Note: The IBM Subsystem Device Driver does not support the Compaq open system in a clustering environment. To have failover protection on an open system, the IBM Subsystem Device Driver requires a minimum of two adapters. You can run the Subsystem Device Driver with one Fibre Channel adapter, but you have no failover protection. The maximum number of adapters supported is 32.

1.6.2 Compaq Tru64 UNIX Version 4.0x host system

In this topic we explain how to install and configure a Tru64 UNIX Version 4.0x host system. Before you configure the ESS volumes, ensure that the system and Fibre Channel adapter cards are at the minimum supported firmware revision.

1.6.3 Verifying the Compaq configuration

The following procedures tell you how to verify the configuration for the host adapter cards and firmware for a Compaq host system:

1. From the console prompt, type **show version**

This is an example of what you see when you type the **show version** command.

```
version V5.8-2, 21-Jul-2000 17:16:08
```

2. Type **#view /var/adm/messages** to display the log file that shows the system version and the version of the Fibre Channel adapter cards.

This is an example of how the firmware version and the host adapter is displayed

```
Sep 21 14:17:10 osplcpq-ds20 vmunix: Firmware revision: 5.8
```

```
Sep 21 14:17:10 osplcpq-ds20 vmunix: KGPSA adapter: Driver Rev 1.09 :  
F/W Rev 2.22X1(1.13) : wwn 1000-0000-c922-d469
```

1.6.4 Operating system device recognition

You can access the disks in the special files. You can find the special files in the **/dev** directory. The disk block special file starts with **rz** and the character disk special file starts with **rrz**. Compaq Tru64 UNIX uses the following syntax to identify the disk special files:

[r] rz [L] [B] [P]

In this syntax:

- ▶ **[L]** — LUN letter. If the LUN is 0, no letter is specified for the LUN. You can specify LUNs 2 - 7 for letters b - h.
- ▶ **[B]** — Bus number multiplied by 8 plus the target number
- ▶ **[P]** — Disk partition from a - h

Following is an example of the output that is displayed from **boot log** or the **scu show edt** command). This shows a disk at LUN 0, bus SCSI 16, target 0.

```
[r]rz[LUN 0][16 * 8 + 0][a-h] = rz128a, rz128b, rz128c, rz128d, rz128e, rz128f,  
rz128g, rz128h,rrz128a, rrz128b, rrz128c, rrz128d, rrz128e, rrz128f, rrz128g,  
rrz128h
```

You do not need to perform any special operations on the Compaq system to view the ESS volumes. Ensure that the host Fibre Channel is already configured. If it is not configured, you must perform the instructions that come with the card to install the driver. Make all the hardware connections, configure the ESS, and restart the host system. If you can see the disks when you restart, all the special files are created automatically.

If you cannot restart the host system, you can configure the devices manually by typing **#scu scan edt**. You can check to see if you can see the volumes by typing **#scu show edt**. To create the device special files type **#doconfig -d**. No kernel rebuild is necessary.

When the ESS is configured, ensure that you can see the disks from the Compaq system. Wait for the system to restart or use the **file** command on the device special files. During restart, the following messages should be displayed:

```
emx0 at pci1 slot 7 KGPSA-BC : Driver Rev 1.21 : F/W Rev 2.22X1(1.13) : wwn
1000-0000-c922-d469
emx0: emx_assign_fcp_id: nport at DID 0x21300 assigned tgt id 10 - out of range
for CAM
scsi16 at emx0 slot 0
rzl128 at scsi16 target 0 lun 0 (LID=0) (IBM 2105F20 1013)
rzb128 at scsi16 target 0 lun 1 (LID=1) (IBM 2105F20 1013)
rzc128 at scsi16 target 0 lun 2 (LID=2) (IBM 2105F20 1013)
rzd128 at scsi16 target 0 lun 3 (LID=3) (IBM 2105F20 1013)
rze128 at scsi16 target 0 lun 4 (LID=4) (IBM 2105F20 1013)
rzf128 at scsi16 target 0 lun 5 (LID=5) (IBM 2105F20 1013)
rzg128 at scsi16 target 0 lun 6 (LID=6) (IBM 2105F20 1013)
rzh128 at scsi16 target 0 lun 7 (LID=7) (IBM 2105F20 1013)
```

When the host system completes the restart, login as the root. Use the **file** command to ensure that all the special files were created.

Use the **file** command to check the major number. If there is a description of the disk, you will not find the model of the disk. Use the information from the output of the previous command or calculate the special file syntax.

This shows an example of what you see when you type **#file /dev/rrz*128c**.

```
/dev/rrz128c: character special (8/262146) SCSI #16 2105F20 disk #1024 (SCSI
ID #0) (SCSI LUN #0)
/dev/rrzb128c: character special (8/262210) SCSI #16 2105F20 disk #1025 (SCSI
ID #0) (SCSI LUN #1)
/dev/rrzc128c: character special (8/262274) SCSI #16 2105F20 disk #1026 (SCSI
ID #0) (SCSI LUN #2)
/dev/rrzd128c: character special (8/262338) SCSI #16 2105F20 disk #1027 (SCSI
ID #0) (SCSI LUN #3)
/dev/rrze128c: character special (8/262402) SCSI #16 2105F20 disk #1028 (SCSI
ID #0) (SCSI LUN #4)
```

```

/dev/rrzf128c: character special (8/262466) SCSI #16 2105F20 disk #1029 (SCSI
ID #0) (SCSI LUN #5)
/dev/rrzg128c: character special (8/262530) SCSI #16 2105F20 disk #1030 (SCSI
ID #0) (SCSI LUN #6)
/dev/rrzh128c: character special (8/262594) SCSI #16 2105F20 disk #1031 (SCSI
ID #0) (SCSI LUN #7)

```

If a special file does not have a disk associated with it, the output of the **#file** command will only show the part concerning the major number. No disk model is displayed.

This is an example showing a disk model:

```

/dev/rrzh128c: character special (8/262594)

```

Type **#disklabel rz128** to see if the disks have a valid disklabel. Then type **#/dev/rrz128a**

The following shows an example of the output from these commands:

```

type: SCSI
disk: 2105F20
label:
flags: dynamic_geometry
bytes/sector: 512
sectors/track: 64
tracks/cylinder: 30
sectors/cylinder: 1920
cylinders: 1017
sectors/unit: 1953152
rpm: 7200
interleave: 1
trackskew: 0
cylinderskew: 0
headswitch: 0          # milliseconds
track-to-track seek: 0 # milliseconds
drivedata: 0
8 partitions:
#      size      offset  fstype  [fsize bsize  cpg] #
a:    131072        0   unused        0    0 # (Cyl.  0 - 68*)
b:    262144   131072   unused        0    0 # (Cyl. 68*- 204*)
c:    1953152        0   unused        0    0 # (Cyl.  0 - 1017*)
d:         0         0   unused        0    0 # (Cyl.  0 - -1)
e:         0         0   unused        0    0 # (Cyl.  0 - -1)
f:         0         0   unused        0    0 # (Cyl.  0 - -1)
g:    819200   393216   unused        0    0 # (Cyl. 204*- 631*)
h:    740736  1212416   unused        0    0 # (Cyl. 631*- 1017*)

```

If there is no disklabel on the disk, type the following command to write the label to the disk:

```
#disklabel -rw rzh128 shark
```

Substitute your specific disk with the **rzh128** command. Always specify the **rz** file without the partition.

1.6.5 Configuring AdvFS

Before you create an AdvFS file system, you must design a structure by assigning a file domain and the file sets. Type the following commands to create an AdvFS file system with file sets:

1. **# cd /**
2. **# mkfdmn -rw /dev/rzXc vol1_dom**
3. **# mkfset vol1_dom vol1**
4. **# mkdir /vol1 # mount vol1_dom#vol1 /vol1**

To display all mounted devices, type: **df -k**

1.6.6 Configuring devices to mount automatically

To enable an AdvFS file system to start automatically, add an entry to the **/etc/fstab** file to the mount command during startup. The following shows an example of a modified **/etc/fstab** file.

The lines that are shown in bold are the lines that were entered since the initial operating system installation:

```
# root_domain#root / advfs rw,userquota,groupquota 0 0
/proc /proc procfs rw 0 0
usr_domain#usr /usr advfs rw,userquota,groupquota 0 0
/dev/rz8b swap1 ufs sw 0 2
vol1_dom#vol1 /vol1 advfs rw,userquota,groupquota 0 2
vol2_dom#vol1 /vol2 advfs rw,userquota,groupquota 0 2
vol3_dom#vol1 /vol3 advfs rw,userquota,groupquota 0 2
vol4_dom#vol1 /vol4 advfs rw,userquota,groupquota 0 2
vol5_dom#vol1 /vol5 advfs rw,userquota,groupquota 0 2
vol6_dom#vol1 /vol6 advfs rw,userquota,groupquota 0 2
vol7_dom#vol1 /vol7 advfs rw,userquota,groupquota 0 2
vol8_dom#vol1 /vol8 advfs rw,userquota,groupquota 0 2
vol9_dom#vol1 /vol9 advfs rw,userquota,groupquota 0 2
vol10_dom#vol1 /vol10 advfs rw,userquota,groupquota 0 2
```

When the host system starts, it mounts all volumes that you created in 1.6.5, “Configuring AdvFS” on page 22.

1.7 Hewlett Packard 9000 FC host system attachment

This topic describes the host system requirements and provides procedures to attach an ESS to a Hewlett Packard 9000 host system with Fibre Channel adapters.

1.7.1 Attachment requirements

This section lists the requirements for attaching the ESS to your host system:

- ▶ Ensure that you have all of the items listed in Equipment requirements.
- ▶ Ensure that you have the installation script files. The script file is on the compact disc that you receive with the ESS.
- ▶ Ensure that you have 1 MB minimum of hard disk space available to install the 2105inst script file.
- ▶ Check the LUN limitations for your host system; see Table 2.
- ▶ Ensure that you have the documentation for your host system and the IBM Enterprise Storage Server User's Guide. The User's Guide is on the compact disc that you receive with the ESS.

See the following Web site for details about the release level for your operating system:

www.storage.ibm.com/hardsoft/products/ess/supserver.htm

Either you or an IBM service support representative (SSR) must perform the following tasks to install and configure an ESS.

1. The IBM SSR installs the ESS by using the procedures in the IBM Enterprise Storage Server Service Guide.
2. Either you or an IBM SSR defines the Fibre Channel host system with the worldwide port name identifiers. For the list of worldwide port names, see 1.12, "Locating the worldwide port name" on page 69.
3. Either you or an IBM SSR defines the Fiber Channel port configuration if you did not do it during the installation of the ESS or Fibre Channel adapters.

Note: Use the information on the logical configuration work sheet in the IBM Enterprise Storage Server Configuration Planner that you previously filled out.

4. Either you or an IBM SSR configures the host system for the ESS by using the instructions in your host system publications.

Note: The IBM Subsystem Device Driver 1.1.3 supports the Hewlett Packard host system in a clustering environment. To have failover protection on an open system, the IBM Subsystem Device Driver requires a minimum of two Fibre Channel adapters. The maximum number of Fibre Channel adapters supported is 16 for a total of 16 Fibre Channel ports.

1.7.2 Installing the 2105 host install script file

This section provides the instructions to install the 2105 host install script file from a compact disc for the ESS on your host system.

Before installing the 2105 host install script file, connect the host system to the ESS. See Attaching an ESS to an open-systems host with SCSI adapters.

Install the 2105 host install script from a compact disc. You must have superuser authority to complete these instructions. The following must be taken into account:

1. You can only install and run the IBM 2105 set queue depth program (version 2.7.1.00) on HP-UX operating system 10.01 or later.
2. If an ESS was used with the SAN Data Gateway on a Hewlett Packard host system and the LUNs were created as Fibre Channel devices by spoofing HP-UX and describing the ESS device types as Sun or other non-Hewlett Packard Fibre Channel device types, the LUNs created under this false ID must be deleted and recreated as Hewlett Packard device types.

Perform the following steps to install the 2105 host install script from a compact disc.

1. If you do not already have a directory called /SD_CDROM, type **mkdir /SD_CDROM** to create a new directory.
2. Insert the compact disc into the CD-ROM drive.
3. Mount the drive as a file system.
4. Type: **ioscan -fnkC disk**
Look for the device name on the list with a name of either compact disc or DVD.
Type: **mount -o cdcase /dev/dsk/c_t_d_ / SD_CDROM**
Replace /dev/dsk/c_t_d_ with the device special file found in step 4.
5. Type: **swinstall -s /SD_CDROM/hp-common/IBMis.depot**

6. From the Software Selection window, click **IBMis_tag**.
7. From the Action menu, click **Mark** for Install.
8. When you see the word Yes next to the IBMis_tag product, go to the Action menu and click **Install**.
9. When the analysis completes with no errors (Status- Ready), click **OK**.
10. Click **Yes** in the Confirmation window to begin the installation.

A window opens, notifying you that the installation is complete and that the system needs to be restarted.
11. Click **OK** to continue.

1.7.3 Configuring the IBM 2105 for clustering

This section describes how to configure a Hewlett Packard host system for clustering.

The steps to configure MC/ServiceGuard with the IBM 2105 are the same as the steps in the Hewlett Packard high availability documentation located at the following Web site:

www.docs.hp.com/hpux/ha/index.html

If you are not familiar with the steps in the Hewlett Packard documentation, IBM recommends that the system administrator attend the Hewlett Packard class about high availability features.

After you configure your host for normal operating system access, the 2105 acts as a normal disk device in the MC/ServiceGuard configuration. IBM recommends that you create volume groups that contain the volumes using the Hewlett Packard logical volume manager. This method of disk management is more reliable, easier, and more flexible to manage than whole-disk management techniques.

Creating volume groups also allows the implementation of PV-Links, Hewlett Packard's built-in multipathing software, for highly available disks such as the IBM 2105. To establish PV-Links, perform the following steps:

1. Create the volume group, using the path to the volumes that you want as the primary path to the data.
2. Extend the volume group with the path to the volumes intended as alternate paths. The logical volume manager reads the label on the disk and knows that it is an alternate path to one of the volumes in the group. The logical volume manager labels the volume. For example, if a host has access to a volume on a 2105 with the device nodes c2t0d0 and c3t0d0, you can use the c2 path as primary and create the volume group using only the c2t0d0 path.

3. Extend the volume group to include the c3t0d0 path. When you issue a **vgdisplay -v** command on the volume group, it lists c3t0d0 as an alternate link to the data.

1.8 Novell NetWare FC host system attachment

This topic describes how to attach a Novell NetWare host system to an IBM ESS with the following adapter cards:

- ▶ QLogic QLA2100F
- ▶ QLogic QLA2200F

Note: The IBM SAN Fibre Channel Switch 2109 S08 and IBM SAN Fibre Channel Switch 2109 S16 are supported for Novell NetWare. The IBM SAN Data Gateway 2108 Model G07 is not supported for Novell NetWare.

1.8.1 Installing the QLogic QLA2100F adapter card

This section tells you how to attach an ESS to a Novell NetWare host system with the QLogic QLA2100F adapter card.

Single-port Fibre Channel interfaces with the QLogic QLA2100F adapter card support the following loop modes:

- ▶ Target
- ▶ Initiator
- ▶ Target and initiator

Note: The arbitrated loop topology is the only topology available for the QLogic QLA2100F adapter card.

Perform the following steps to install the QLogic QLA2100F adapter card:

Note: The following steps are an example of a configuration. The configuration for your adapter might differ.

1. Install the QLogic QLA2100F adapter card in the host system.
2. Connect the cable to the ESS port.
3. Restart the server.
4. Press Alt+Q to get to the FAST!Util menu.
5. From the Configuration Settings menu, click **Host Adapter Settings**.

6. From the Advanced Adapter Settings menu, press the Down Arrow to highlight LUNs per target. Press Enter.
7. Press the Down Arrow to find and highlight **256**. Press Enter.
8. Press Esc.
9. To save the changes, click **Yes**. Press Enter.
10. Restart the server.

1.8.2 Installing the QLogic QLA2200F adapter card

This section tells you how to attach an ESS to a Novell NetWare host system with the QLogic QLA2200F adapter card. Single-and dual port Fibre Channel interfaces with the QLogic QLA2200F adapter card support the following public and private loop modes:

- ▶ Target
- ▶ Public initiator
- ▶ Private initiator
- ▶ Target and public initiator
- ▶ Target and private initiator

Perform the following steps to install the QLogic QLA2200F adapter card:

Note: The following steps are an example of a configuration. The configuration for your adapter might differ.

1. Install the QLogic QLA2200F adapter card in the host system.
2. Connect the cable to the ESS port.
3. Restart the server.
4. Press Alt+Q to get to the FAST!Util Command panel.
5. From the Configuration Settings menu, click Host Adapter Settings and set the parameters and values from the Host Adapter Settings menu as follows:
 - a. Host adapter BIOS: **Disabled**
 - b. Frame size: **2048**
 - c. Loop reset delay: **5 (minimum)**
 - d. Adapter hard loop ID: **Disabled**
6. From the Advanced Adapter Settings menu, press the Down Arrow key to highlight LUNs per target; then press Enter. Set the parameters and values from the Advanced Adapter Settings menu as follows:

- a. Execution throttle: **240**
 - b. Fast command posting: **Enabled**
 - c. >4 GB addressing: **Disabled for 32 bit systems**
 - d. LUNs per target: **0**
 - e. Enable LIP reset: **No**
 - f. Enable LIP full login: **No**
 - g. Enable target reset: **Yes**
 - h. Login retry count: **20 (minimum)**
 - i. Port down retry count: **20 (minimum)**
 - j. Driver load RISC code: **Enabled**
 - k. Enable database updates: **No**
 - l. Disable database load: **No**
 - m. IOCB allocation: **256**
 - n. Extended error logging: **Disabled (might be enabled for debugging)**
7. Press Esc to return to the Configuration Settings menu.
 8. From the Configuration Settings menu, scroll down to **Extended Firmware Settings**. Press Enter.
 9. From the Extended Firmware Settings menu, scroll down to **Connection Options** to open the Option and Type of Connection window.
 10. Press Enter.
 11. Select from one of the following options:
 - **0**: Loop only
 - **1**: Point-to-point
 - **2**: Loop preferred (If you cannot use arbitrated loop, then default to point-to-point.)
 - **3**: Point-to point, otherwise loop (If you cannot use point-to-point, default to arbitrated loop.)
- Note:** If you connect the ESS directly to the host system, the option you select must match the port connections on the ESS.
12. Press Esc.
 13. To save the changes, click **Yes**. Press Enter.
 14. Restart the server.

1.8.3 Loading the current Fibre Channel adapter driver

Perform the following steps to load the current driver onto the QLogic adapter card.

1. Go to the following Web site:
www.qlogic.com
2. From the home page, click **Driver Download**.
3. Click **Drivers**.
4. Click **Fibre-Channel Adapter Drivers**.
5. Click **QLA2xxx drivers**.
6. Click **Novell NetWare**.
7. Click **QLogic Vx.xxx** where *V* is the version and *x.xxx* is the version level of the file name.
8. In the Save As window, find the current driver file, *xxxxxxx.exe*, where *xxxxxxx* is the driver file name.

Note: We recommend that you save the file to a floppy diskette.

9. Click **Save**.
10. Close the Web site.
11. From your host system Start menu, click **Run**.
12. In the Run window, ensure that the drive letter in the field is the same as the drive letter where you saved the *xxxxxxx.exe* file. If no drive letter appears, type the letter of the drive where you saved the driver file.
13. Type the driver file name after *x:*, where *x* is the drive letter you specified to save the file.
14. Type the directory name where you want to put the file. Click **Zip**.
15. Click **OK** to unzip the current driver file.

1.8.4 Installing the Fibre Channel adapter drivers

Perform the following steps to install the Fibre Channel adapter drivers.

1. From the NetWare server console, type **nwconfig** for NetWare 5.0 or 5.1 or **load install** for NetWare 4.x.
2. Select **Driver Options**.
3. Select **Configure Disk and Storage Device Drivers**.
4. In the SCSI Adapters window, click the **Drivers** tab.

5. Click **“Select an additional driver.”**
6. Press the Insert key.
7. Insert a floppy diskette with the QLogic drivers into the A:\ drive of the NetWare server. Press Enter.
The available driver is displayed.
8. Select the driver for the QLogic card and press Enter.
9. Select **Modify driver parameters** and enter the slot number of the QLogic card into the slot number parameter.
10. Set the **Scan All Luns** parameter to **Yes**.
11. Press Tab and select Save Parameters and Load Driver.
12. Exit the **nwconfig** or install utility.
13. If storage has already been assigned to the server from the ESS Specialist, type **SCAN FOR NEW DEVICES**, **SCAN ALL LUNS**, and **LIST DEVICES**.

The ESS hosted volumes are displayed in the devices list. Create volumes using the **nwconfig** utility if necessary.

1.8.5 ESS QLogic QLA2100F/2200F host adapter card

To configure the host adapter card, use the IBM TotalStorage ESS Specialist.

1.9 Sun FC host system attachment

This topic tells you how to change the Sun system kernel. Before you start, you must meet several attachment requirements.

1.9.1 Attachment requirements

This section lists the requirements for attaching the ESS to your host system:

- ▶ Ensure that you have all the equipment you need.
- ▶ Ensure that there are enough Fibre Channel adapters installed in the server to handle the total LUNS you want to attach. The maximum LUNs per adapter in a Fibre Channel arbitrated loop topology is 50.
- ▶ Ensure that you have the documentation for your host system, and that the *IBM Enterprise Storage Server User's Guide* is on the compact disc that you receive with the ESS.

- Solaris 2.6, Solaris 7, and Solaris 8 require patches to ensure that the host and the ESS function correctly. See Table 1-6 for the minimum revision level that is required for each Solaris patch ID.

Table 1-6 Solaris 2.6, 7, and 8 minimum revision level patches for Fibre Channel

Solaris 2.6	Solaris 7	Solaris 8
105181-23 kernel update	106541-12 kernel update	108528-03 kernel update
105356-16 sd, ssd drivers	106924-06 isp driver	109524-02 ssd driver
105580-16 glm driver	106925-04 glm driver	109657-01 isp driver
105600-19 isp driver	107147-08 pci driver	108974-03 sd, uata drivers
Not applicable	107458-10 dad, sd, ssd, uata drivers	Not applicable

- Review the Sun host SCSI adapter device driver installation and configuration utility documents for additional Solaris patches that you might need.
- See the following Web site for details about the release level for your operating system:

www.storage.ibm.com/hardsoft/products/ess/supserver.htm

Either you or an IBM service support representative (SSR) must perform the following tasks to install and configure an ESS.

1. The IBM SSR installs the ESS by using the procedures in the IBM Enterprise Storage Server Service Guide
2. Either you or an IBM SSR defines the Fibre Channel host system with the worldwide port name identifiers.
3. Either you or an IBM SSR defines the Fiber Channel port configuration if you did not do it during the installation of the ESS or Fibre Channel adapters.

Note: Use the information on the logical configuration work sheet in the IBM Enterprise Storage Server Configuration Planner you should have previously filled out.

4. Either you or an IBM SSR configures the host system for the ESS by using the instructions in your host system publications.

Note: The IBM Subsystem Device Driver does not support the Sun host system in a clustering environment. To have failover protection on an open system, the IBM Subsystem Device Driver requires a minimum of two Fibre Channel adapters. The maximum number of Fibre Channel adapters supported is 16 for a total of 16 Fibre Channel ports.

The two IBM Subsystem Device Drivers for Sun host systems support:

- ▶ Sun host hardware platforms limited to 32-bit mode and all Sun host systems running Solaris 2.6.
- ▶ Sun host hardware platforms with 64-bit mode capabilities running Solaris 7 or Solaris 8.

1.9.2 Installing the Emulex LP8000 adapter card

This section tells you how to attach an ESS to a Windows NT host system with the Emulex LP8000 adapter card. Single-and dual-port Fibre Channel interfaces with the Emulex LP8000 adapter card support the following public and private loop modes:

- ▶ Target
- ▶ Public initiator
- ▶ Private initiator
- ▶ Target and public initiator
- ▶ Target and private initiator

The ESS supports increased connectivity with the use of Fibre Channel (SCSI-FCP and FICON) directors. Specific details on status, availability, and configuration options for the Fibre Channel directors supported by the ESS are available on the Web site:

www.storage.ibm.com/hardsoft/products/ess/supserver.htm

Perform the following steps to install the adapter card:

1. Turn off and unplug the computer.
2. Remove the computer case.
3. Remove the blank panel from an empty PCI bus slot.
4. Insert the host adapter board into the empty PCI bus slot. Press firmly until seated.
5. Secure the mounting bracket for the adapter to the case with the panel screw.

6. Replace the computer case by tightening the screws on the case or use the clamp to secure the cover.

1.9.3 Downloading the current Emulex Fibre Channel adapter driver

This section tells you how to install the Emulex driver:

1. Plug in and restart your host system.
2. Go to the following Web site:
`www.emulex.com`
3. From the left navigation menu for Quick Links, click **Documentation, Drivers and Software**.
4. From the Fibre Channel menu, select the adapter model.
5. Click **Drivers for Solaris**.
6. Click SCSI/IP v4.xxx where xxx equals the level of the driver for Solaris.
7. Click **Download Now for an Intel driver** or **Download Now for a SPARC driver**.
8. From the File Download window, click **Save this file to disk**.
9. Click **OK**.
10. In the Save As window, click **Save**.
A window that shows the progress of the download is displayed.
11. When the download completes, click **Close**.

Note: If you downloaded the driver file to a host system other than a Sun, you must transfer the file to the Sun host system. Otherwise, go to "Installing the Emulex LP8000 Fibre Channel adapter drivers".

1.9.4 Installing the Fibre Channel adapter drivers

Perform the following steps to install the Fibre Channel adapter drivers:

1. Login as root.
2. Type **emlxtmp** to create a temporary directory.
3. Type **cd emlxtmp** to change the directory to the temporary directory. If you are downloading the file from the FTP site or have the file on the CD-ROM, go to step 5.
4. Type **/etc/init.d/volmgt stop** and then **umount /dev/fd** to copy the tar file from a floppy diskette
5. Copy or download the device driver file to the temporary directory.

6. If the file is in the format of filename.tar.Z, type `uncompress filename.tar.Z`.
7. Type `tar xvf lpfc-sparc.tar` to “untar” the drive file from the temporary directory.
8. Type `pkgadd -d pwd` to install the package.

Note: An installation script displays that prompts you to answer a number of questions. For each question, enter the appropriate response. Or, press Enter to each question to accept the default setting.

9. Specify the package number or press Enter to accept all packages.
10. Type `y` or `n` to answer the prompt that reads: **Rebuild manual pages database for section 7d [y,n?]**:. If you typed `y`, go to step 11.

Note: Rebuilding the manual pages can take up to ten minutes. If you do not want to build the manual pages, type `n`. You can run the command later.

11. At the prompt that reads Use IP networking over Fibre Channel [y,n?]: type `y` or `n`. If you typed `y`, go to step 12.
12. Type the name of the network host name for the adapter.

Note: The network host name identifies the host adapter on a Fibre Channel network and is associated with a unique IP address.

13. Edit the `/etc/hosts` file to add the IP address to the host name.

Note: If you have more than one adapter in the system, you must create a `hostname.lpf#file` for each adapter.

14. At the prompt that reads:

Do you want to continue with the installation of <lpfc>

Type `y` to proceed with the installation, or type `n` to undo all the settings and end the installation.

15. At the prompt that reads:

Select package(s) you wish to process (or 'all' to process all packages).
(default:all) [?,??,q]:

Type `q`.

Note: We recommend that you configure the host adapter parameters before you shut down and restart the host system.

16. At the system prompt, type **shutdown** to restart the host system.
17. Login as root.

18. Update the parameter list. See Table 1-8 on page 42.
19. Restart the host system.

Note: You must restart your system twice if IP is enabled.

1.9.5 Installing the JNI PCI adapter card

This section tells you how to attach an ESS to a Sun host system with the JNI PCI adapter card:

1. Perform the following steps to install the JNI PCI adapter card.
2. Turn off and unplug the computer.
3. Remove the computer case.
4. Remove the blank panel from an empty PCI bus slot.
5. Insert the host adapter board into the empty PCI bus slot. Press firmly until seated.
6. Secure the mounting bracket for the adapter to the case with the panel screw.
7. Replace the computer case by tightening the screws on the case or use the clamp to secure the cover.

1.9.6 Downloading the current JNI PCI Fibre Channel adapter driver

This section tells you how to download the JNI PCI Fibre Channel adapter driver:

1. Plug in and restart your host system.
2. Go to the following Web site:
www.jni.com
3. From the navigation menu at the top of the page, click **Drivers**.
4. From the menu for Locate Driver by Product, click **FCI-1063**.
5. From the menu for FCI-1063, find the section for Solaris -JNI. Click **fca-pci.pkg**.
6. In the dialog box for File Download, click **Save this file to disk**. Click **OK**.
7. In the Save As dialog box, create a temporary folder.
8. Click **Save**.
9. When you click **Save**, a window displays that shows the progress of the download.
10. When the download completes, click **Close**.

11. If you downloaded the driver file from a Sun host system, go to **Installing the JNI PCI Fibre Channel adapter driver**. If you downloaded the driver file from a non-Sun host system, transfer the drive file to a Sun host system.

1.9.7 Installing the JNI PCI Fibre Channel adapter driver

Perform the following steps to install the JNI PCI Fibre Channel adapter drivers:

1. Go to the following Web site:
`www.jni.com`
2. From the navigation menu at the top of the page, click **Drivers**.
3. From the menu for Locate Driver by Product, click **FCI-1063**.
4. From the menu for FCI-1063, find the section for Solaris -JNI. Click **readme.txt**.
5. Print the readme.txt file.
6. Follow the instructions in the readme.txt file to install the JNI PCI adapter card.
7. Update the parameter list. See Table 1-7 on page 42.
8. Restart the host system.

1.9.8 Installing the JNI SBUS adapter card

This section tells you how to attach an ESS to a Sun host system with the JNI SBUS adapter card.

Perform the following steps to install the JNI SBUS adapter card:

1. Turn off and unplug the computer.
2. Remove the computer case.
3. Remove the blank panel from an empty SBUS slot.
4. Insert the host adapter board into the empty SBUS slot. Press firmly until seated.
5. Secure the mounting bracket for the adapter to the case with the panel screw.
6. Replace the computer case by tightening the screws on the case or use the clamp to secure the cover.

1.9.9 Downloading current JNI SBUS Fibre Channel adapter driver

This section tells you how to download the JNI SBUS Fibre Channel adapter driver:

1. Plug in and restart your host system.
2. Go to the following Web site:
`www.jni.com`
3. From the navigation menu at the top of the page, click **Drivers**.
4. From the menu for Locate Driver by Product, click **FC64-1063**.
5. From the menu for FCI-1063, find the section for Solaris -JNI. Click **fcw.pkg**.
6. In the dialog box for File Download, click **Save** this file to disk. Click **OK**.
7. In the Save As dialog box, create a temporary folder.
8. Click **Save**.
9. When you click **Save**, a window displays that shows the progress of the download.
10. When the download completes, click **Close**.
11. If you downloaded the driver file from a Sun host system, go to **Installing the JNI SBUS Fibre Channel adapter driver**. If you downloaded the driver file from a non-Sun host system, transfer the drive file to a Sun host system.

1.9.10 Installing the JNI SBUS Fibre Channel adapter driver

Perform the following steps to install the JNI SBUS Fibre Channel adapter drivers.

1. Go to the following Web site:
`www.jni.com`
2. From the navigation menu at the top of the page, click **Drivers**.
3. From the menu for Locate Driver by Product, click **FC64-1063**.
4. From the menu for FC64-1063, find the section for Solaris -JNI. Click **readme.txt**.
5. Print the readme.txt file.
6. Follow the instructions in the readme.txt file to install the JNI SBUS adapter card.
7. Update the parameter list. See Table 1-7 on page 42.
8. Restart the host system.

1.9.11 Installing the QLogic QLA2200F adapter card

This section tells you how to attach an ESS to a host system with the QLogic QLA2200F adapter card. Single-and dual-port Fibre Channel interfaces with the QLogic QLA2200F adapter card support the following public and private loop modes:

- ▶ Target
- ▶ Public initiator
- ▶ Private initiator
- ▶ Target and public initiator
- ▶ Target and private initiator

Perform the following steps to install the QLogic QLA2200F adapter card:

1. Install the QLogic QLA2200F adapter card in the host system.
2. Connect the cable to the ESS port.
3. Restart the server.
4. Press Alt+Q to get to the FAST!Util menu.
5. From the Configuration Settings menu, select **Host Adapter Settings**.
6. From the Host Adapter Settings menu, set the following parameters and values:
 - a. Host adapter BIOS: **Disabled**
 - b. Frame size: **2048**
 - c. Loop reset delay: **5** (minimum)
 - d. Adapter hard loop ID: **Disabled**
7. From the Advanced Adapter Settings menu, press the Down Arrow to highlight LUNs per target. Press Enter. Set the parameters and values from the Advanced Adapter Settings menu as follows:
 - a. Execution throttle: **100**
 - b. Fast command posting: **Enabled**
 - c. >4 GB addressing: **Disabled for 32 bit systems**
 - d. LUNs per target: **0**
 - e. Enable LIP reset: **No**
 - f. Enable LIP full login: **No**

Note: In a clustering environment, set Enable LIP full login to Yes.

- g. Enable target reset: **Yes**
- h. Login retrycount: **20** (minimum)
- i. Port down retrycount: **20** (minimum)
- j. Driver load RISC code: **Enabled**
- k. Enable database updates: **No**

- l. Disable database load: **No**
- m. IOCB allocation: **256**
- n. Extended error logging: **Disabled** (might be enabled for debugging)

Note: The Enable LIP reset, Enable LIP full logon, and Enable target reset parameters control the behavior of the adapter when Windows NT tries to do a SCSI bus reset. You must perform a target reset to make cluster failovers work. Use the SCSI bus device reset option to clear SCSI reservations.

- 8. Press Esc to return to the Configuration Settings menu.
- 9. From the Configuration Settings menu, scroll down to **Extended Firmware Settings**. Press Enter.
- 10. From the Extended Firmware Settings menu, scroll down to **Connection Options** to open the Option and Type of Connection window.
- 11. Select one of the following options:
 - a. **0**: Loop only
 - b. **1**: Point-to-point (preferred setting)
 - c. **2**: Loop preferred (If you cannot use arbitrated loop, then default to point-to-point)
 - d. **3**: Point-to-point, otherwise loop (If you cannot use point-to-point, default to arbitrated loop).

Note: If you connect the ESS directly to the host system, the option you select must match the port connections on the ESS. However, if you connect through a switch, the options do not need to match the port connections because the ESS is point-to-point. The appropriate HBA on the server must also support point-to-point connection on a direct connection. Currently, disparate vendors do not function properly in a direct point-to-point connection. This statement is not true if you connect through a switch because the ESS is point-to-point.

- 12. Press Esc.
- 13. Save the changes. Highlight **Yes**.
- 14. Restart the server.

1.9.12 Downloading the current Fibre Channel adapter driver

Perform the following steps to load the current driver onto the QLogic adapter card.

- 1. Go to the following Web site:
www.qlogic.com

2. From the home page, click **Driver Download**.
3. Click **Use QLogic Drivers**.
4. Click **Fibre Channel Adapter Drivers and Software**.
5. In table, click **QLA22xx drivers**.
6. From the menu for Software and Drivers available for, click **Solaris**.
7. From the table for QLA22xx Driver Download Page, Solaris Sparc 2.6/2.7/2.8, and Current released (Sparc) driver PCI to FC Adapter, click **Link to Driver**.
8. Download the file and save as required.

1.9.13 Installing the Fibre Channel adapter drivers

Perform the following steps to install the Fibre Channel adapter drivers.

Note: If you are installing the Fibre Channel adapter for the first time, you must specify the correct topology. You must also select the appropriate device mapping driver.

1. Go to the following Web site:
www.qlogic.com
2. From the home page, click **Driver Download**.
3. Click **Use Qlogic Drivers button**.
4. Click **Fibre Channel Adapter Drivers and Software**.
5. In the table for Qlogic Fibre Channel Adapters, click **QLA22xx**.
6. From the menu for Software and Drivers available for, click **Solaris**.
7. From the table for QLA22xx Driver Download Page, Solaris Sparc 2.6/2.7/2.8, and Current released (Sparc) driver PCI to FC Adapter, click **Read Me**.
 This action displays the contents of the Read Me file, which contains the instructions to install the driver file.
8. Update the parameter file if necessary. See Table 1-9 on page 43.

1.9.14 Configuring host device drivers

The following procedures tell you how to update the Solaris SCSI driver configuration file to enable access to target and LUN pairs configured on the ESS.

1. Change to the directory by typing: **cd /kernel/drv**
2. Back up the sd.conf file in this subdirectory.
3. Edit the sd.conf file to add support for the target and LUN pairs that are configured on the host system. The following is an example of the lines that you would add to the file to access LUNs 0-7 on target 8.

```
name="sd" class="scsi"  
target=0 lun=0;  
name="sd" class="scsi"  
target=0 lun=1;  
name="sd" class="scsi"  
target=0 lun=2;  
name="sd" class="scsi"  
target=0 lun=3;  
name="sd" class="scsi"  
target=0 lun=4;  
name="sd" class="scsi"  
target=0 lun=5;  
name="sd" class="scsi"  
target=0 lun=6;  
name="sd" class="scsi"  
target=0 lun=7;
```

Note: Do not add duplicate target and LUN pairs.

Type **reboot ---r** from the Open Windows window to shut down and restart the Sun host system with the kernel reconfiguration option, or type **boot ---r** from the OK prompt after you shut down.

The Fibre Channel adapters that are supported for Sun attachment to the ESS are capable of full fabric support. IBM recommends that all Fibre Channel driver configurations include worldwide port name, worldwide node name, port ID, or host-bus-adapter binding of target LUN pairs.

Binding of target LUN pairs implements the Solaris SCSI driver configuration file or in the Fibre Channel host-bus-adapter configuration file installed by the adapter software package. Refer to the vendor adapter documentation and utilities for detailed configuration instructions.

1.9.15 Tuning HBA configuration files

Fibre Channel host-bus-adapter configuration files are tunable for host system reliability and performance.

See Table 1-7 for the recommended configuration settings for the host-bus-adapter for a JNI FC64-1063 and a JNI FCI-1063.

Table 1-7 Recommended JNI HBA settings

Parameters	Recommended settings
fca_nport	0: Default. Initializes on a loop 1: Recommended for fabric. Initializes as an N_Port.
public loop	0: Default. Recommended. Initializes according to which fca_nport is set as disabled.
ip_disable	0: Default. IP side of the driver is enabled. 1: Recommended for fabric. IP side of the adapters is completely disabled.
failover	See the following Web site for details on how to set the parameter setting: www.storage.ibm.com/hardsoft/products/ess/supserver.htm
busy_retry_delay	500: Recommended. Delay between retries after device returns a busy response for a command.
scsi_probe_delay	5000: Recommended. Delay before SCSI probes are allowed during startup.

See Table 1-8 for recommended host bus adapter configuration file parameters for an Emulex LP8000 adapter

Table 1-8 Recommended Emulex LP8000 HBA settings.

Parameters	Recommended settings
automap	1: Default. Automatically assigns SCSI IDs to FCP targets.
fcp-on	1: Default. Turn on FCP.
lun-queue-depth	16: Recommended when there are less than 17 LUNs per adapter. Set value = 256/(Total LUNs per adapter) when there are more than 16 LUNs per adapter. If your configuration includes more than one LP8000 adapter per server, calculate the lun-queue-depth value using the adapter with the most LUNs attached.
no-device-delay	15: Recommended. Delay to failback and I/O.
network-on	0: Default. Recommended for fabric. Do not turn on IP networking 1: Turn on IP networking.

Parameters	Recommended settings
scan-down	2: Recommended. Use inverted ALPA map and cause target assignment in private loop.
topology	2: Recommended for fabric. Point-to-point mode only. 4: Recommended for non-fabric. Arbitrated loop mode only.
zone-rscn	0: Default 1: Recommended for fabric; check name server for RSCNs.

In Table 1-9 we show the recommended configuration file parameters for the host bus adapters for the QLogic QLA2200F adapter.

Table 1-9 Recommended QLA2200F HBA settings

Parameters	Recommended settings
hba0-max-frame-length	2048
hba0-max-iocb-allocation	256
hba0-execution-throttle	31
hba0-login-timeout	4
hba0-login-retry-count	1
hba0-fabric-retry-count	10
hba0-enable-adapter-hard-loop	0
hba0-adapter-hard-loop-l	0
hba0-enable-64bit-addressing	0
hba0-enable-LIP-reset	0
hba0-enable-LIP-full-login	1
hba0-enable-target-reset	0: Non-clustered 1: Clustered
hba0-reset-delay	5
hba0-port-down-retry-count	30
hba0-link-down-error	1
hba0-loop-down-timeout	60
hba0-connection-options	1: Fabric connection 2: Direct connection
hba0-device-configuration-mode	1
hba0-fc-tape	0
hba0-command-completion-option	1

1.9.16 Installing the IBM Subsystem Device Driver

The following instructions explain how to install the IBM Subsystem Device Driver from a compact disc. You can use the IBM Subsystem Device Driver in conjunction with the IBM Copy Services command-line interface program.

1. Type **ps -ef |grep vold** to ensure that the volume manager is running. This command displays the `/usr/sbin/vold` process. If it does not display, type:
/etc/init.d/volmgt start
2. Insert the IBM Subsystem Device Driver CD into the CD-ROM drive. A File Manager window opens showing the paths for the Subsystem Device Driver package subdirectories.

Note: You must be on the host console to see this window.

3. Change to the subdirectory that contains the Subsystem Device Driver package.
 - a. For Sun host hardware platforms limited to 32-bit mode and for all Sun host systems running Solaris 2.6,type:
cd /cdrom/unnamed_cdrom/Sun32bit
 - b. For Sun host hardware platforms with 64-bit capabilities running Solaris 7 or Solaris 8,type:
cd /cdrom/unnamed_cdrom/Sun64bit
4. Type **pkgadd -d** to initiate the Package Add menu.
5. Select the option number for the IBM DPO driver (IBMdpo), and press Enter.
6. Select **y** to continue the installation for all prompts until the package installation is complete.
7. Select **q** and press Enter to exit the package options menu.
8. Type **cd** to change back to the root directory.
9. Type **eject cdrom** and press Enter to remove the Subsystem Device Driver CD.

Edit the **profile** file in the root directory and add the lines shown below to include the IBM DPO subdirectory in the system path:

```
PATH=$PATH:/opt/IBMdpo/bin
export PATH
```

The IBM DPO driver subdirectory will be automatically added to the path when the system is restarted.

1.9.17 Setting the Sun host system parameters

The following sections contain the procedures to set the Sun host system parameters for optimum performance on the ESS with the following adapters:

- ▶ JNI
- ▶ Emulex
- ▶ QLogic

JNI adapters

The following sections contain the procedures to set the Sun host system parameters for optimum performance on the ESS with the JNI adapter:

1. Type **cd /etc** to change to the /etc subdirectory.
2. Back up the system file in the subdirectory.
3. Edit the system file and set the following parameters for servers with configurations that use JNI adapters:

sd_max_throttle

This `sd_max_throttle` parameter specifies the maximum number of commands that the `sd` driver can queue to the host bus adapter driver. The default value is 256, but you must set the parameter to a value less than or equal to a maximum queue depth for each LUN connected. Determine the value using the following formula:

$256/(\text{LUNs per adapter})$

The parameter shows thirty two 2105 LUNs attached to controller 1, (c1t#d#), and forty eight 2105 LUNs attached to controller 2, (c2t#d#). The value for `sd_max_throttle` is calculated using the controller with the highest number of LUNs attached. The `sd_max_throttle` parameter for the ESS LUNs in this example would be set by adding the following line to the /etc/system file:

```
set sd:sd_max_throttle=5
```

sd_io_time

This parameter specifies the timeout value for disk operations. Add the following lines to the /etc/system file to set the `sd_io_time` parameter for the ESS LUNs:

```
set sd:sd:_io_time=0x78
```

sd_retry_count

This parameter specifies the retry count for disk operations. Add the following line to the /etc/system file to set the `sd_retry_count` parameter for the ESS LUNs:

```
set sd:sd_retry_count=5
```

maxphys

This parameter specifies the maximum number of bytes you can transfer for each SCSI transaction. The default value is 126976 (124 KB). If the I/O block size that you requested exceeds the default value, the request is broken into more than one request. The value should be tuned to the intended use and application requirements. For maximum bandwidth set the maxphys parameter by adding the following line to the /etc/system file:

```
set maxphys=8388608
```

If you are using Veritas volume manager on the ESS LUNs, you must set the VxVM max I/O size parameter (vol_maxio) to match the maxphys parameter. If you set the maxphys parameter to 8388608, add the following line to the /etc/system file to also set the VxVM I/O size to 8 MB:

```
set vxio:vol_maxio=16384
```

Emulex or QLogic adapters

The following sections contain the procedures to set the Sun host system parameters for optimum performance on the ESS with the Emulex adapter:

1. Type `cd /etc` to change to the /etc subdirectory.
2. Back up the system file in the subdirectory.
3. Edit the system file and set the following parameters for servers with configurations that only use Emulex adapters.

sd_io_time

This parameter specifies the timeout value for disk operations. Add the following line to the /etc/system file to set the sd_io_time parameter for the ESS LUNs:

```
set sd:sd:_io_time=0x78
```

sd_retry_count

This parameter specifies the retrycount for disk operations. Add the following line to the /etc/system file to set the sd_retry_count parameter for the ESS LUNs:

```
set sd:sd_retry_count=5
```

maxphys

This parameter specifies the maximum number of bytes you can transfer for each SCSI transaction. The default value is 126976 (124 KB). If the I/O block size that you requested exceeds the default value, the request is broken into more than one request. The value should be tuned to the intended use and application requirements. For maximum bandwidth, set the maxphys parameter by adding the following line to the /etc/system file:

```
set maxphys=8388608
```

If you are using Veritas volume manager on the ESS LUNs, you must set the VxVM max I/O size parameter (vol_maxio) to match the maxphys parameter. If you set the maxphys parameter to 8388608, add the following line to the /etc/system file to also set the VxVM I/O size to 8 MB:

```
set vxio:vol_maxio=16384
```

1.10 Windows NT 4.0 FC host system attachment

In this topic we tell you how to attach an ESS to a Windows NT host system with the following Fibre Channel adapters.

- ▶ QLogic QLA2100F adapter card
- ▶ QLogic QLA2200F adapter card
- ▶ Emulex LP8000 adapter card

This topic also tells you how to install, download, and configure the adapter cards.

1.10.1 Attachment requirements

This section lists the requirements for attaching the ESS to your host system:

- ▶ Check the LUN limitations for your host system.
- ▶ Ensure that you have the documentation for your host system and the *IBM Enterprise Storage Server User's Guide* (the *User's Guide* is on the compact disc that you receive with the ESS).
- ▶ See the following Web site for details about the release level for your operating system:

www.storage.ibm.com/hardsoft/products/ess/supserver.htm

Either you or an IBM SSR must perform the following tasks to install and configure an ESS.

1. The IBM SSR installs the ESS by using the procedures in the IBM Enterprise Storage Server Service Guide
2. Either you or an IBM SSR defines the Fibre Channel host system with the worldwide port name identifiers.

3. Either you or an IBM SSR defines the Fiber Channel port configuration if you did not do it during the installation of the ESS or Fibre Channel adapters.

Note: Use the information on the logical configuration work sheet in the IBM Enterprise Storage Server Configuration Planner that you should have previously filled out.

You or an IBM SSR configures the host system for the ESS by using the instructions in your host system publications.

Notes:

1. The IBM Subsystem Device Driver 1.2.1 supports the Windows NT 4.0 host system in a clustering environment. To have failover protection on an open system, the IBM Subsystem Device Driver requires a minimum of two Fibre Channel adapters. The maximum number of Fibre Channel adapters supported is 16 for a total of 16 Fibre Channel ports.
2. To improve performance, IBM recommends that you map to the LUNs for the target volumes of the Windows NT host until you need access to the data on the target volume. Perform the LUN mapping after the PPRC operation and immediately before you need access to the data. You must restart the host system before you can access the data on the target volume. You can greatly reduce the time it takes for the host system to restart if you perform the LUN mapping. Otherwise, the time to reboot could take 10 minutes per PPRC target volume.

1.10.2 Installing the QLogic QLA2100F adapter card

This section tells you how to attach an ESS to a Windows NT host system with the QLogic QLA2100F adapter card. Single-port Fibre Channel interfaces with the QLogic QLA2100F adapter card support the following loop modes:

- ▶ Target
- ▶ Initiator
- ▶ Target and initiator

Notes:

1. The arbitrated loop topology is the only topology available for the QLogic QLA2100F adapter card.
2. When you install the QLA2100F adapter card on your host system, IBM recommends that you install the card on a host system with four processors. This ensures that you do not have a problem when you start the host system.

Perform the following steps to install the QLogic QLA2100F adapter card:

1. Install the QLogic QLA2100F adapter card in the host system.
2. Connect the cable to the ESS port.
3. Restart the host system.
4. Press Alt+Q to get to the FAST!Util menu.
5. From the Configuration Settings menu, select **Host Adapter Settings**.
6. From the Advanced Adapter Settings menu, press the Down Arrow to highlight LUNs per target; then press Enter.
7. Use the Down Arrow to find and highlight **256**. Press Enter.
8. Press Esc.
9. To save the changes, click **Yes**. Press Enter.
10. Restart the server.

1.10.3 Installing the QLogic QLA2200F adapter card

This section tells you how to attach an ESS to a Windows NT host system with the QLogic QLA2200F adapter card. Single-and dual-port Fibre Channel interfaces with the QLogic QLA2200F adapter card support the following public and private loop modes:

- ▶ Target
- ▶ Public initiator
- ▶ Private initiator
- ▶ Target and public initiator
- ▶ Target and private initiator

Perform the following steps to install the QLogic QLA2200F adapter card:

1. Install the QLogic QLA2200F adapter card in the host system.
2. Connect the cable to the ESS port.
3. Restart the server.
4. Press Alt+Q to get to the FAST!Util menu.
5. From the Configuration Settings menu, select **Host Adapter Settings**.
6. From the Host Adapter Settings menu, set the following parameters and values:
 - a. Host adapter BIOS: **Disabled**
 - b. Frame size: **2048**
 - c. Loop reset delay: **5** (minimum)
 - d. Adapter hard loop ID: **Disabled**

7. From the Advanced Adapter Settings menu, press the Down Arrow to highlight LUNs per target. Press Enter. Set the parameters and values from the Advanced Adapter Settings menu as follows:
 - a. Execution throttle: **100**
 - b. Fast command posting: **Enabled**
 - c. >4 GB addressing: **Disabled** for 32 bit systems
 - d. LUNs per target: **0**
 - e. Enable LIP reset: **No**
 - f. Enable LIP full login: **No**

Note: In a clustering environment, set Enable LIP full login to Yes.

- g. Enable target reset: **Yes**
- h. Login retrycount: **20** (minimum)
- i. Port down retrycount: **20** (minimum)
- j. Driver load RISC code: **Enabled**
- k. Enable database updates: **No**
- l. Disable database load: **No**
- m. IOCB allocation: **256**
- n. Extended error logging: **Disabled** (might be enabled for debugging)

Note: The Enable LIP reset, Enable LIP full logon, and Enable target reset parameters control the behavior of the adapter when Windows NT tries to do a SCSI bus reset. You must perform a target reset to make cluster failovers work. Use the SCSI bus device reset option to clear SCSI reservations.

8. Press Esc to return to the Configuration Settings menu.
9. From the Configuration Settings menu, scroll down to **Extended Firmware Settings**. Press Enter.
10. From the Extended Firmware Settings menu, scroll down to **Connection Options** to open the Option and Type of Connection window.
11. Select one of the following options:
 - a. **0**: Loop only
 - b. **1**: Point-to-point (preferred setting)
 - c. **2**: Loop preferred (If you cannot use arbitrated loop, then default to point-to-point)
 - d. **3**: Point-to point, otherwise loop (If you cannot use point-to-point, default to arbitrated loop).

Note: If you connect the ESS directly to the host system, the option you select must match the port connections on the ESS. However, if you connect through a switch, the options do not need to match the port connections because the ESS is point-to-point. The appropriate HBA on the server must also support point-to-point connection on a direct connection. Currently, disparate vendors do not function properly in a direct point-to-point connection. This statement is not true if you connect through a switch because the ESS is point-to-point.

12. Press Esc.
13. Save the changes. Highlight **Yes**.
14. Restart the server.

1.10.4 Downloading the current Fibre Channel adapter driver

Perform the following steps to load the current driver onto the QLogic adapter card:

1. Go to the following Web site:
www.qlogic.com
2. From the home page, click **Driver Download**.
3. Click **Use QLogic Drivers**.
4. Click **IBM Enterprise Subsystems Division approved drivers**.
5. Click **IBM Approved QLA22xx drivers** or **IBM Approved QLA21xx drivers**.
6. Click **Link to Driver for Windows NT**.
7. In the window for File Download, click **Save this Program to Disk**.
8. Insert a floppy diskette into your floppy disk drive.
9. In the field for Save in, change the drive letter to **A:**
10. Create a folder called Temp in the Save As window.
11. Click **Save**.
12. When the download completes, click **Open Folder**.
13. From the A:\ window, double-click the icon for the driver file you downloaded to the floppy diskette.
14. In the window for WinZip Self-Extractor, type **a:**
15. Click **Unzip**.

When the unzip process completes, you should see a message that says, x files unzipped successfully, where x equals the number of files you unzipped.

16. Click **OK**.
17. Click **Close** to close the window for WinZip Self-Extractor.

1.10.5 Installing the Fibre Channel adapter drivers

Perform the following steps to install the Fibre Channel adapter drivers.

Note: If you are installing the Fibre Channel adapter for the first time, you must specify the correct topology. You must also select the appropriate device mapping driver.

1. From your Windows NT desktop, double-click the icon for **My Computer**.
2. Double-click the icon for **Control Panel**.
3. Double-click the icon for **SCSI Adapters**.
4. In the SCSI Adapters window, click the **Drivers** tab.
5. Click **Add**.
6. In the Install Drivers window, click **Have Disk**.
7. In the Install from Disk window, ensure that the drive letter in the field for Copy Manufacturer's Files From is the drive letter you specified to save the .exe file to when you downloaded the driver previously.
8. Type the name of the current driver file in the field for Copy Manufacturer's Files From after the drive letter prompt.
9. Click **OK**.
10. Click **Cancel** to exit.
11. Restart your host system.

1.10.6 QLogic QLA2100F or QLogic QLA2200F configuration

To configure the host adapter card, use the IBM TotalStorage ESS Specialist.

1.10.7 Configuring for availability and recoverability

This section describes how to ensure optimum availability and recoverability when you attach an IBM ESS to a Windows NT host system. You must set the time-out value associated with the supported host bus adapters to 240 seconds. The setting is consistent with the configuration for IBM SSA adapters and disk subsystems when attached to a Windows NT host system.

The host bus adapter uses the time-out parameter to bound its recovery actions and responses to the disk subsystem. The value exists in different places in the system configuration. You can retrieve and use it in different ways, depending on the type of host bus adapter. The following instructions tell you how to modify the value safely in either the Windows NT registry or in the device adapter parameters.

1.10.8 Setting the TimeOutValue registry

The following instructions tell you how to set the time-out value registry:

1. From the Run menu or command prompt, type:
Regedt32.exe
2. Navigate to the following registry key:
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Disk
3. Look for the value called TimeOutValue. If the value called TimeOutValue does not exist, go to step 3a. If the TimeOutValue exists, go to step 4.
 - a. Click **Edit ->Add Value...**
 - b. For ValueName, click **TimeOutValue**.
 - c. For data type, click **REG_DWORD**.
 - d. Click **OK**.
 - e. For data, click **f0**.
 - f. For radix, click **Hex**.
 - g. Click **OK**.
4. If the value exists and is less than 0x000000f0 (240 decimal), perform the following steps to increase it to 0xf0:
 - a. Click **TimeOutValue**.
 - b. Click **Edit ->DWORD...**
 - c. For data, click **f0**.
 - d. For radix, click **hex**.
 - e. Click **OK**.
5. Exit the Regedt32 program.
6. Restart your Windows NT server for the changes to take effect.

1.10.9 Installing Emulex LP8000 adapter cards

This section tells you how to attach an ESS to a Windows NT host system with Emulex LP8000 adapter cards. Single-and dual-port Fibre Channel interfaces with the Emulex LP8000 adapter cards support the following public and private loop modes:

- ▶ Target
- ▶ Public initiator
- ▶ Private initiator
- ▶ Target and public initiator
- ▶ Target and private initiator

The ESS supports increased connectivity with the use of Fibre Channel (SCSI-FCP and FICON) directors. Specific details on status, availability, and configuration options for the Fibre Channel directors supported by the ESS are available on the Web:

www.storage.ibm.com/hardsoft/products/ess/supserver.htm

Perform the following steps to install the Emulex LP8000 adapter card:

1. Turn off and unplug the computer.
2. Remove the computer case.
3. Remove the blank menu from an empty PCI bus slot.
4. Insert the host adapter board into the empty PCI bus slot. Press firmly until seated.
5. Secure the mounting bracket for the adapter to the case with the panel screw.
6. Replace the computer case by tightening the screws on the case or use the clamp to secure the cover.

1.10.10 Downloading the current Fibre Channel adapter driver

This section tells you how to install the port driver:

1. Plug in and restart your host system.
2. Go to the following Web site:
www.emulex.com
3. From the Quick Links menu, click **Documentation, Drivers, and Software**.
4. Click the **host adapter type** from the host adapter menu.
5. Click **Drivers for Windows NT**.
6. Click **Specialized Drivers**.
7. Click **Fibre Channel Port xxxxx** where **xxxxx** is the name of the adapter driver.
8. Click the **Download Now** button.

9. From the File Download window, click the appropriate radio button and proceed as indicated:
 - To open this file from its current location, go to step 9.
 - To save this file to disk, go to step 16.
 10. In the window for Winzip, click **I agree**.
 11. In the window for WinZip Wizard - Welcome, click **Next**.
 12. In the window for WinZip Wizard - Select Zip File xxxxxxxx.zip, where xxxxxxxx is the name of the file, highlight the file that you want to unzip.
 13. Click **Next**.
 14. In the window for WinZip Wizard - Unzip, click **Unzip now**.

A progress indicator window appears when the operation to unzip the file completes. A window opens to display the following file names:

 - Lpscsi
 - Lputilnt
 - Oemsetup
 - Readme
 - Txtsetup.oem
 15. Double-click **Readme** to get the instructions to install the Fibre Channel adapter driver. Print the Readme file.
 16. In the window for WinZip Wizard -Unzip Complete, click **Close**.
 17. Ensure that the name of the file you want to download is displayed in the window.
 18. Click **Save** to download and unzip the file to your hard drive.
- A window opens that indicates the progress of the download operation. When the progress indicator window closes, the download is complete.

1.10.11 Installing the Fibre Channel adapter drivers

Perform the following steps to install the Fibre Channel adapter drivers.

Note: If you are installing the Fibre Channel adapter for the first time, you must specify the correct topology. You must also select the appropriate device mapping driver.

1. From your desktop, click **Start Settings**.
2. Double-click **Control Panel**.
3. Double-click **SCSI Adapters**.
4. Click the **Drivers** tab.

5. Click **Add** to create a list of drivers.
A window opens that indicates the progress. When the window closes, you should see a window called Install Driver.
6. From the Install Driver window, click **Have Disk**.
7. Enter the path to the driver file that you downloaded and click **OK**.
8. To install the driver, highlight the line that lists the driver you want and click **OK**.

Note: The driver affects every adapter in the system. If you have more than one adapter that requires different parameter settings, you must change the parameter settings with the port utility and restart your host system.

9. Click **Yes** to restart the host system.
10. After you restart your host system, click **Start Settings**.
11. Double-click **Control Panel**.
12. Double-click **SCSI Adapters**.
13. Click the **Drivers** tab and verify that the Emulex SCSI driver is present and started.
14. Click the **Devices** tab and verify that the host adapter is on the list.

1.10.12 Parameter settings for the Emulex LP8000 on Windows NT

In Table 1-10 we show the recommended host bus adapter configuration file parameters for an Emulex LP8000 adapter. The settings are for an IBM ESS model F20 that is attached through a switch using the fabric, automap SCSI devices port driver, and an IBM ESS model F20 that is attached directly, using arbitrated loop, automap SCSI devices port driver.

Table 1-10 Recommended LP8000 settings for NT

Parameters	Recommended settings
Automatically map SCSI devices	Checked (enabled)
Query name server for all N-ports	Checked (enabled)
Allow multiple paths to SCSI targets	Checked (enabled)
Point-to-point	Not checked (disabled) for direct attach Not shown for the fabric attach
Register for state change	Checked (enabled)
Use report LUNs	Checked (enabled)

Parameters	Recommended settings
Use name server after RSCN	Checked (enabled) if fabric attached using soft zoning Not checked (disabled)
Lun mapping	Checked (enabled)
Automatic lun mapping	Checked (enabled)
Scan in device ID order	Not checked (disabled)
Enable class 2 for SCSI devices	Not checked (disabled)
Report unknown SCSI devices	Not checked (disabled)
Look for disappearing devices	Not checked (disabled)
Translate queue full to busy	Not checked (disabled)
Use bus reset status for retries	Not checked (disabled)
Retry unit attention	Not checked (disabled)
Retry PLOGI open failures	Not checked (disabled)
Maximum Number of LUNs	Equal to or greater than the number of the ESS LUNs available to the host bus adapter
Maximum Queue Depth	8
Link Timer	30 seconds
Retries	64
E_D_TOV	2000 milliseconds
AL_TOV	15 milliseconds
Wait ready timer	45 seconds
Retry timer	2000 milliseconds
R_A_TOV	2 seconds
ARB_TOV	1000 milliseconds
Link Control	Not shown for direct attach
Topology	Point-to-point (fabric) Arbitrated loop (direct connect)
Link speed	Auto

1.10.13 Verifying the configuration

The following tells you how to determine whether or not your Windows NT 4.0 host system is configured for storage:

1. From the Windows NT desktop, right-click **Start**.
2. Partition new Fibre Channel drives with Microsoft Disk Administrator. See your Windows NT documentation for more instructions.
3. Click **Explore** and verify that you can see the Fibre Channel drives.
4. Select a large file (for example, a 9 MB file), and copy it to a Fibre Channel drive.

1.10.14 Verifying networking

The following tells you how to determine whether or not your Windows NT 4.0 host system is configured for networking:

1. From the Windows NT desktop, click **Start Programs Command Prompt**.
2. From the command line, ping a known TCP/IP address.

1.11 Windows 2000 FC host system attachment

This topic tells you how to attach an ESS to a Windows 2000 host system with the following Fibre Channel adapters.

- ▶ QLogic QLA2100F adapter card
- ▶ QLogic QLA2200F adapter card
- ▶ Emulex LP8000 adapter card

1.11.1 Attachment requirements

This section lists the requirements for attaching the ESS to your host system:

- ▶ Check the LUN limitations for your host system.
- ▶ Ensure that you have the documentation for your host system and the *IBM Enterprise Storage Server User's Guide* (the *User's Guide* is on the compact disc that you receive with the ESS).

See the following Web site for details about the release level for your operating system:

www.storage.ibm.com/hardsoft/products/ess/supserver.htm

Either you or an IBM SSR must perform the following tasks to install and configure an ESS:

1. The IBM SSR installs the ESS by using the procedures in the IBM Enterprise Storage Server Service Guide
2. You or an IBM SSR defines the Fibre Channel host system with the worldwide port name identifiers.
3. You or an IBM SSR defines the Fiber Channel port configuration if you did not do it during the installation of the ESS or Fibre Channel adapters. For the list of worldwide port names, see 1.12, “Locating the worldwide port name” on page 69.

Note: Use the information on the logical configuration work sheet in the IBM Enterprise Storage Server Configuration Planner that you should have previously filled out.

4. Either you or an IBM SSR configures the host system for the ESS by using the instructions in your host system publications.

Note:

1. The IBM Subsystem Device Driver 1.2.1 does not support the Windows 2000 host system in a clustering environment. To have failover protection on an open system, the IBM Subsystem Device Driver requires a minimum of two Fibre Channel adapters. The maximum number of Fibre Channel adapters supported is 16 for a total of 16 Fibre Channel ports.
2. To improve performance, IBM recommends that you map to the LUNs for the target volumes of the Windows 2000 host until you need access to the data on the target volume. Perform the LUN mapping after the PPRC operation and immediately before you need access to the data. You must restart the host system before you can access the data on the target volume. You can greatly reduce the time it takes for the host system to restart if you perform the LUN mapping. Otherwise, the time to restart could take 10 minutes per PPRC target volume.

1.11.2 Installing the QLogic QLA2100F adapter card

This section tells you how to attach an ESS to a Windows NT host system with the QLogic QLA2100F adapter card. Single-port Fibre Channel interfaces with the QLogic QLA2100F adapter card support the following loop modes:

- ▶ Target
- ▶ Initiator
- ▶ Target and initiator

Note:

1. The arbitrated loop topology is the only topology available for the QLogic QLA2100F adapter card.
2. When you install the QLA2100F adapter card on your host system, IBM recommends that you install the card on a host system with four processors. This ensures that you do not have a problem when you start the host system.

Perform the following steps to install the QLogic QLA2100F adapter card:

1. Install the QLogic QLA2100F adapter card in the host system.
2. Connect the cable to the ESS port.
3. Restart the host system.
4. Press Alt+Q to get to the FAST!Util menu.
5. From the Configuration Settings menu, select **Host Adapter Settings**.
6. From the Advanced Adapter Settings menu, press the Down Arrow to highlight LUNs per target and then press Enter.
7. Use the Down Arrow to find and highlight **256**. Press Enter.
8. Press Esc.
9. To save the changes, click **Yes**. Press Enter.
10. Restart the server.

1.11.3 Installing the QLogic QLA2200F adapter card

This section tells you how to attach an ESS to a Windows 2000 host system with the QLogic QLA2200F adapter card. Single-and dual-port Fibre Channel interfaces with the QLogic QLA2200F adapter card support the following public and private loop modes:

- ▶ Target
- ▶ Public initiator
- ▶ Private initiator
- ▶ Target and public initiator
- ▶ Target and private initiator

Perform the following steps to install the QLogic QLA2200F adapter card:

1. Install the QLogic QLA2200F adapter card in the host system.
2. Connect the cable to the ESS port.
3. Restart the server.

4. Press Alt+Q to get to the FAST!Util menu.
5. From the Configuration Settings menu, select **Host Adapter Settings**.
6. From the Host Adapter Settings menu, set the following parameters and values:
 - a. Host adapter BIOS: **Disabled**
 - b. Frame size: **2048**
 - c. Loop reset delay: **5** (minimum)
 - d. Adapter hard loop ID: **Disabled**
7. From the Advanced Adapter Settings menu, press the Down Arrow to highlight LUNs per target. Press Enter. Set the parameters and values from the Advanced Adapter Settings menu as follows:
 - a. Execution throttle: **100**
 - b. Fast command posting: **Enabled**
 - c. >4 GB addressing: **Disabled** for 32 bit systems
 - d. LUNs per target: **0**
 - e. Enable LIP reset: **No**
 - f. Enable LIP full login: **No**

Note: In a clustering environment, set Enable LIP full login to **Yes**.

- g. Enable target reset: **Yes**
- h. Login retrycount: **20** (minimum)
- i. Port down retrycount: **20** ((minimum)
- j. Driver load RISC code: **Enabled**
- k. Enable database updates: **No**
- l. Disable database load: **No**
- m. IOCB allocation: **256**
- n. Extended error logging: **Disabled** (might be enabled for debugging)

Note: The Enable LIP reset, Enable LIP full logon, and Enable target reset parameters control the behavior of the adapter when Windows NT tries to do a SCSI bus reset. You must perform a target reset to make cluster failovers work. Use the SCSI bus device reset option to clear SCSI reservations.

8. Press Esc to return to the Configuration Settings menu.
9. From the Configuration Settings menu, scroll down to **Extended Firmware Settings**. Press Enter.

10. From the Extended Firmware Settings menu, scroll down to **Connection Options** to open the Option and Type of Connection window.
 11. Select one of the following options:
 - a. **0**: Loop only
 - b. **1**: Point-to-point (preferred setting)
 - c. **2**: Loop preferred (If you cannot use arbitrated loop, then default to point-to-point)
 - d. **3**: Point-to point, otherwise loop (If you cannot use point-to-point, default to arbitrated loop).
- Note:** If you connect the ESS directly to the host system, the option you select must match the port connections on the ESS. However, if you connect through a switch, the options do not need to match the port connections because the ESS is point-to-point. The appropriate HBA on the server must also support point-to-point connection on a direct connection. Currently, disparate vendors do not function properly in a direct point-to-point connection. This statement is not true if you connect through a switch because the ESS is point-to-point.
12. Press Esc.
 13. Save the changes. Highlight **Yes**.
 14. Restart the server.

1.11.4 Downloading the current Fibre Channel adapter driver

Perform the following steps to load the current driver onto the QLogic adapter card.

1. Go to the following Web site:
www.qlogic.com
2. From the home page, click **Driver Download**.
3. Click **Use QLogic Drivers**.
4. Click **IBM Enterprise Subsystems Division approved drivers**.
5. Click **IBM Approved QLA22xx drivers or IBM Approved QLA21xx drivers**.
6. Click **Link to Drivers**.
7. In the window for File Download, click **Save this Program to Disk**.
8. Insert a floppy diskette into your floppy disk drive.
9. In the field for Save in, change the drive letter to **A:**
10. Create a folder called Temp in the Save As window.

11. Click **Save**.
12. When the download completes, click **Open Folder**.
13. From the A:\ window, double-click the icon for the driver file you downloaded to the floppy diskette.
14. In the window for WinZip Self-Extractor, type **a:**
15. Click **Unzip**.
When the unzip process completes, you should see a message that says, *x* files unzipped successfully, where *x* equals the number of files you unzipped.
16. Click **OK**.
17. Click **Close** to close the window for WinZip Self-Extractor.

1.11.5 Installing the Fibre Channel adapter drivers

Perform the following steps to install the Fibre Channel adapter drivers.

Note: If you are installing the Fibre Channel adapter for the first time, you must specify the correct topology. You must also select the appropriate device mapping driver.

1. From your Windows 2000 desktop, double-click the icon for **My Computer**.
2. Double-click the icon for **Control Panel**.
3. Double-click the icon for **SCSI Adapters**.
4. In the SCSI Adapters window, click the **Drivers** tab.
5. Click **Add**.
6. In the Install Drivers window, click **Have Disk**.
7. In the Install from Disk window, ensure that the drive letter in the field for Copy Manufacturer's Files From is the drive letter you specified to save the .exe file to when you downloaded the driver previously.
8. Type the name of the current driver file in the field for Copy Manufacturer's Files From after the drive letter prompt.
9. Click **OK**.
10. Click **Cancel** to exit.
11. Restart your host system.

1.11.6 QLogic QLA2100F or QLogic QLA2200F configuration

To configure the host adapter card, use the IBM TotalStorage ESS Specialist.

1.11.7 Configuring for availability and recoverability

This section describes how to ensure optimum availability and recoverability when you attach an IBM ESS to a Windows 2000 host system. You must set the time-out value associated with the supported host bus adapters to 240 seconds. The setting is consistent with the configuration for IBM SSA adapters and disk subsystems when attached to Windows 2000 host system.

The host bus adapter uses the time-out parameter to bound its recovery actions and responses to the disk subsystem. The value exists in different places in the system configuration. You can retrieve and use it in different ways depending on the type of host bus adapter. The following instructions tell you how to modify the value safely in either the Windows 2000 registry or in the device adapter parameters.

1.11.8 Setting the TimeOutValue registry

The following instructions tell you how to set the time-out value registry:

1. From the Run menu or command prompt, type:
Regedt32.exe
2. Navigate to the following registry key:
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Disk
3. Look for the value called TimeOutValue. If the value called TimeOutValue does not exist, go to step 3a. If the TimeOutValue exists, go to step 4.
 - a. Click **Edit ->Add Value...**
 - b. For ValueName, click **TimeOutValue**.
 - c. For data type, click **REG_DWORD**.
 - d. Click **OK**.
 - e. For data, click **f0**.
 - f. For radix, click **Hex**.
 - g. Click **OK**.
4. If the value exists and is less than 0x000000f0 (240 decimal), perform the following steps to increase it to 0xf0:
 - a. Click **TimeOutValue**.
 - b. Click **Edit ->DWORD...**
 - c. For data, click **f0**.
 - d. For radix, click **hex**.
 - e. Click **OK**.
5. Exit the Regedt32 program.
6. Restart your Windows 2000 server for the changes to take effect.

1.11.9 Installing Emulex LP8000 adapter cards

This section tells you how to attach an ESS to a Windows 2000 host system with Emulex LP8000 adapter cards. Single-and dual-port Fibre Channel interfaces with the Emulex LP8000 adapter cards support the following public and private loop modes:

- ▶ Target
- ▶ Public initiator
- ▶ Private initiator
- ▶ Target and public initiator
- ▶ Target and private initiator

The ESS supports increased connectivity with the use of Fibre Channel (SCSI-FCP and FICON) directors. Specific details on status, availability, and configuration options for the Fibre Channel directors supported by the ESS are available on the Web:

www.storage.ibm.com/hardsoft/products/ess/supserver.htm

Perform the following steps to install the Emulex LP8000 adapter card:

1. Turn off and unplug the computer.
2. Remove the computer case.
3. Remove the blank menu from an empty PCI bus slot.
4. Insert the host adapter board into the empty PCI bus slot. Press firmly until seated.
5. Secure the mounting bracket for the adapter to the case with the panel screw.
6. Replace the computer case by tightening the screws on the case, or use the clamp to secure the cover.

1.11.10 Downloading the current Fibre Channel adapter driver

This section tells you how to install the port driver.

1. Plug in and restart your host system.
2. Go to the following Web site:
www.emulex.com
3. From the Quick Links menu, click **Documentation, Drivers, and Software**.
4. Click the **host adapter type** from the host adapter menu.
5. Click **Drivers for Windows 2000**.
6. Click **Fibre Channel Port v5-1.27a5**.
7. Click the **Download Now** button.

8. From the File Download window, click the appropriate radio button and proceed as indicated:
 - To open this file from its current location, go to step 9.
 - To save this file to disk, go to step 16.
9. In the window for Winzip, click **I agree**.
10. In the window for WinZip Wizard - Welcome, click **Next**.
11. In the window for WinZip Wizard - Select Zip File xxxxxxxx.zip where xxxxxxxx is the name of the file, highlight the file that you want to unzip.
12. Click **Next**.
13. In the window for WinZip Wizard - Unzip, click **Unzip now**.

A progress indicator window appears when the operation to unzip the file completes. A window opens to display the following file names:

 - Lpscsi
 - Lputilnt
 - Oemsetup
 - Readme
 - Txtsetup.oem
14. Double-click **Readme** to get the instructions to install the Fibre Channel adapter driver. Print the Readme file.
15. In the window for WinZip Wizard -Unzip Complete, click **Close**.
16. Ensure that the name of the file you want to download is displayed in the window.
17. Click **Save** to download and unzip the file to your hard drive.

A window opens that indicates the progress of the download operation. When the progress indicator window closes, the download is complete.

1.11.11 Installing the Fibre Channel adapter drivers

Perform the following steps to install the Fibre Channel adapter drivers.

Note: If you are installing the Fibre Channel adapter for the first time, you must specify the correct topology. You must also select the appropriate device mapping driver.

1. From your desktop, click **Start Settings**.
2. Double-click **Control Panel**.
3. Double-click **SCSIAdapters**.
4. Click the **Drivers** tab.

5. Click **Add** to create a list of drivers.
A window opens that indicates the progress. When the window closes, you should see a window called Install Driver.
6. From the Install Driver window, click **Have Disk**.
7. Enter the path to the driver file that you downloaded and click **OK**.
8. To install the driver, highlight the line that lists the driver you want and click **OK**.

Note: The driver affects every adapter in the system. If you have more than one adapter that requires different parameter settings, you must change the parameter settings with the port utility and restart your host system.

9. Click **Yes** to restart the host system.
10. After you restart your host system, click **Start Settings**.
11. Double-click **Control Panel**.
12. Double-click **SCSI Adapters**.
13. Click the **Drivers** tab and verify that the Emulex SCSI driver is present and started.
14. Click the **Devices** tab and verify that the host adapter is on the list.

1.11.12 Parameter settings for the Emulex LP8000 on Windows 2000

In Table 1-11 we show the recommended host bus adapter configuration file parameters for an Emulex LP8000 adapter. The settings are for an IBM ESS model F20 that is attached through a switch using the fabric, automap SCSI devices port driver, and an IBM ESS model F20 that is attached directly, using arbitrated loop, automap SCSI devices port driver.

Table 1-11 Recommended LP8000 settings for Windows 2000

Parameters	Recommended settings
Automatically map SCSI devices	Checked (enabled)
Query name server for all N-ports	Checked (enabled)
Allow multiple paths to SCSI targets	Checked (enabled)
Point-to-point	Not checked (disabled) for direct attach Not shown for the fabric attach
Register for state change	Checked (enabled)
Use report LUNs	Checked (enabled)

Parameters	Recommended settings
Use name server after RSCN	Checked (enabled) if fabric attached using soft zoning Not checked (disabled)
Lun mapping	Checked (enabled)
Automatic lun mapping	Checked (enabled)
Scan in device ID order	Not checked (disabled)
Enable class 2 for SCSI devices	Not checked (disabled)
Report unknown SCSI devices	Not checked (disabled)
Look for disappearing devices	Not checked (disabled)
Translate queue full to busy	Not checked (disabled)
Use bus reset status for retries	Not checked (disabled)
Retry unit attention	Not checked (disabled)
Retry PLOGI open failures	Not checked (disabled)
Maximum Number of LUNs	Equal to or greater than the number of the ESS LUNs available to the host bus adapter
Maximum Queue Depth	8
Link Timer	30 seconds
Retries	64
E_D_TOV	2000 milliseconds
AL_TOV	15 milliseconds
Wait ready timer	45 seconds
Retry timer	2000 milliseconds
R_A_TOV	2 seconds
ARB_TOV	1000 milliseconds
Link Control	Not shown for direct attach
Topology	Point-to-point (fabric) Arbitrated loop (direct connect)
Link speed	Auto

1.11.13 Verifying the configuration

The following tells you how to determine whether or not your Windows 2000 host system is configured for storage.

1. From the Windows 2000 desktop, right-click **Start**.
2. Partition new Fibre Channel drives with Microsoft Disk Administrator. See your Windows 2000 documentation for more instructions.
3. Click **Explore** and verify that you can see the Fibre Channel drives.
4. Select a large file (for example, a 9 MB file), and copy it to a Fibre Channel drive.

1.11.14 Verifying networking

The following tells you how to determine whether or not your Windows 2000 host system is configured for networking:

1. From the Windows 2000 desktop, click **Start Programs Command Prompt**.
2. From the command line, ping a known TCP/IP address.

1.12 Locating the worldwide port name

This topic tells you how to locate the WWPN value for a host system Fibre Channel adapter for the following operating systems:

- ▶ AS/400 iSeries
- ▶ NUMA-Q and xSeries
- ▶ RS/6000 and pSeries
- ▶ Compaq
- ▶ Hewlett Packard 9000
- ▶ Novell NetWare
- ▶ Sun
- ▶ Windows NT
- ▶ Windows 2000

1.12.1 Fibre Channel port name identification

The WWPN consists of exactly 16 hexadecimal characters (0-9 and A-F). It is used by the ESS to uniquely identify the Fibre Channel adapter card that is installed in your host system. The ESS automatically finds the WWPN for your host Fibre Channel adapter when you attach your host system to the ESS.

If your host system uses more than one Fibre Channel adapter to connect to your ESS, you must add multiple entries to the host system list for this host, one for each Fibre Channel adapter. Each adapter will have its own unique WWPN.

The format and content of the Fibre Channel port identifier are determined by the manufacturer of the link control facility for the applicable Fibre Channel port. The identifier is an eight-byte field, which the Fibre Channel protocols use to uniquely identify the Fibre Channel port. An example is the ESS Fibre Channel port.

You can manually assign a unique worldwide port name for the ESS by performing the steps in the following sections.

1.12.2 Locating the WWPN for an AS/400 iSeries

To locate the WWPN for a AS/400 host system, perform the following steps:

1. On the screen for the AS/400 Main Menu panel, type **strsst**.
2. On the Start Service Tools (STRSST) Sign On panel, type your service tools userid and service tools password.
3. On the System Service Tools (SST) panel, type **1** to select **Start a service tool**.
4. On the Start a Service Tool panel, type **7** to select **Hardware service manager**.
5. On the Hardware Service Manager panel, type **1** to select **Packaging hardware resources (systems, frames, cards, ...)**.
6. On the Packaging Hardware Resources panel, type **9** to select the **System Expansion unit**.
7. On the Packaging Hardware Resources panel, type **8** to select **Multiple Function IOA**.
8. On the Logical Resources Associated with a Packaging Resource panel, type **5** to select **Multiple Function IOA**.
9. On the Auxiliary Storage Hardware Resource Detail panel, locate the field name for Port worldwide name. Note the number in the right hand column.

Note: If you have exchanged a 2766 Fibre Channel IOA in the iSeries system, the IBM 2105 ESS disk unit subsystem must be updated to use the worldwide port name of the new 2766 IOA. You can find the name in the port worldwide name field on the iSeries system by displaying detail on the 2766 IOA Logical Hardware Resource information in Hardware Service Manager in SST/DST.

1.12.3 Locating the WWPN for an IBM eServer or IBM NUMA-Q host

To locate the WWPN for a NUMA-Q host system with an IOC-0210-54 adapter, perform the following steps from the IBM Enterprise Storage Specialist:

1. From the Enterprise Storage Specialist Welcome panel, click **Storage Allocation**.
2. From the Storage Allocation graphical view panel, click **Open System Storage**.
3. From the Open System Storage panel, click **Modify Host Systems**.
4. In the Host Nickname field, type the nickname.
5. In the Host Name field, click either **IBM NUMA Server (WinNt)** or **IBM NUMA Server (UNIX)** from the list.
6. Click the **Down Arrow** to the right of the Host Attachment field.
7. From the list, highlight and click **Fibre Channel Attached**.
8. In the Hostname/IP Address field, type the hostname.
9. Click the **Down Arrow** to the right of the World-Wide Port Name field.
10. Select the worldwide port name from the list.
11. Click **Perform Configuration Update**.

1.12.4 Locating the WWPN for an RS/6000 and pSeries host system

To locate the WWPN for an RS/6000 or pSeries host system, perform the following steps:

1. Login as root.
2. Type `lscfg -v1 fcsx`, where x is the adapter number.

The network address is the Fibre Channel adapter port WWPN value.

Note: The `lscfg -v1 fcsx` ROS level identifies the Fibre Channel adapter firmware level.

1.12.5 Locating the WWPN for a Compaq host system

To locate the WWPN for a Compaq host system, perform the following steps:

1. From the console prompt, type `P0>>>wwidmgr -show ada`.

Following is an example of what is displayed:

```
Probing timeout
item adapter WWN Cur.Topo Next Topo
[0 ]pga0.0.0.7.1 1000-0000-c922-d469 FABRIC FABRIC
[1 ]pgb0.0.0.8.1 2000-0000-c922-6a63 FABRIC FABRIC
[9999 ]All of the above.
```

If you receive the following error:

wwidmgr available only prior to booting. Reinit system and try again.

Type **P00>>>init** and repeat the **wwidmgr** command.

If you receive the following error:

wwidmgr: No such command

If you see this message, type **P00>>set mode diag** and type the **wwidmgr** command again.

If the system is already running, you can find the WWPN in the log file **/var/adm/messages**.

2. Type **#fgrep wwn /var/adm/messages**

Note: The worldwide port names are recorded at each system startup.

This is an example of typing **var/adm/messages**.

```
Nov 9 09:01:16 osplcpq-ds20 vmunix:KGPSA-BC :Driver Rev 1.21 :F/W Rev
2.22X1(1.13):wwn 1000-0000-c922-d469
Nov 10 10:07:12 osplcpq-ds20 vmunix:KGPSA-BC :Driver Rev 1.21 :F/W Rev
2.22X1(1.13):wwn 1000-0000-c922-d469
Nov 13 17:25:28 osplcpq-ds20 vmunix:KGPSA-BC :Driver Rev 1.21 :F/W Rev
2.22X1(1.13):wwn 1000-0000-c922-d469
Nov 14 11:08:16 osplcpq-ds20 vmunix:KGPSA-BC :Driver Rev 1.21 :F/W Rev
2.22X1(1.13):wwn 1000-0000-c922-d469
Nov 15 10:49:31 osplcpq-ds20 vmunix:KGPSA-BC :Driver Rev 1.21 :F/W Rev
2.22X1(1.13):wwn 1000-0000-c922-d469
```

1.12.6 Locating the WWPN for a Hewlett Packard host system

To locate the WWPN for a Hewlett Packard host system, perform the following steps:

1. Go to the root directory.
2. Type **ioscan -fn | more**
3. Look under the description for the Fibre Channel Mass Storage adapter.

For example, to look for the device path name /dev/td1, type:

```
fcmsutil /dev/td1 where /dev/td1
```

1.12.7 Locating the WWPN for a Novell NetWare host system

To locate the WWPN for a Novell NetWare host system with a QLogic adapter, perform the following steps:

1. Restart the server.
2. Press Alt+Q to get the FAST!Util menu.
If you have more than one Fibre Channel adapter installed, you will see a screen that displays all the Fibre Channel adapters. Scroll down to the adapter you want.
3. Press Enter.
4. From the Fast Util!menu, scroll down and select **Select Host Adapter**.
5. Scroll up and highlight **Configuration Settings**. Press Enter.
6. From the Configuration Settings menu, click **Host Adapter Settings**.
7. Write down the host adapter name, for example: 200000E08B00C2D5.

1.12.8 Locating the WWPN for a Sun host system

To locate the WWPN for a Sun host system for a JNI Fibre Channel adapter card, perform the following steps:

1. Go to the root directory
2. Type `dmesg | more`

All the devices in the directory are displayed. When you type `dmesg | more`, you might find the worldwide port name of other Fibre Channel adapters if the driver displays the worldwide port name during a restart.

Scan the output message file for `fcaw` or `jnic`, and the 64-bit number if you have a different JNI Fibre Channel driver. Then look for the worldwide port name, for example, 20000E06940028.,

Note:

1. There is no method to locate the WWPN for a QLogic QLA2100F adapter installed on a Sun host system.
2. If you have multiple host adapters installed, you will see more than one WWPN.

1.12.9 Locating the WWPN for a Windows NT host system

To locate the WWPN for a Windows NT host system with a QLogic adapter, perform the following steps:

1. Restart the server.
2. Press Alt+Q to get the FAST!Util menu.
If you have more than one Fibre Channel adapter installed, a panel displays all the Fibre Channel adapters. Scroll down to the adapter you want. Press Enter.
3. From the Fast!Util menu, scroll down and select **Select Host Adapter**.
4. Scroll up and highlight **Configuration Settings**.
5. Press Enter.
6. From the Configuration Settings menu, click **Host Adapter Settings**.
7. Write down the host adapter name, for example, 200000E08B00C2D5.

1.12.10 Locating the WWPN for a Windows 2000 host system

To locate the WWPN for a Windows 2000 host system with a QLogic adapter, perform the following steps:

1. Restart the server.
2. Press Alt+Q to get the FAST!Util menu.
If you have more than one Fibre Channel adapter installed, all the Fibre Channel adapters are displayed.
3. Scroll down to the adapter you want.
4. Press Enter.
5. From the Fast!Util menu, scroll down and select **Select Host Adapter**.
6. Scroll up and highlight **Configuration Settings**.
7. Press Enter.
8. From the Configuration Settings menu, click **Host Adapter Settings**.
9. Write down the host adapter name, for example: 200000E08B00C2D5.



IBM TotalStorage Enterprise Storage Server configuration

One benefit of a SAN is to implement disk pooling. To do this successfully, we need an easy-to-manage storage server with a Fibre Channel attachment. This is the IBM TotalStorage Enterprise Storage Server, 2105-F20, shown in Figure 2-1.

First we introduce the ESS Web Interface, including the functions to maintain, configure, and administer the ESS. Then we perform storage configuration.

We describe the configuration for an ESS using RAID 5 for use in a Fibre Channel environment. These functions are described in a task-oriented manner. We start with an unconfigured ESS and explain the configuration steps as we progress through our configuration. The final goal is to gain access to the volumes from Fibre Channel attached hosts.

In greater detail, we cover the configuration of the ESS disk groups, host Fibre Channel ports (host FC ports), ESS Fibre Channel ports (ESS FC ports), and the ESS logical volumes to be accessed from the host FC ports. We also modify volume assignments and make changes to the host FC port definition.

For further information on the configuration of the ESS, especially for ESCON and SCSI attachment and the ESS Web Copy Services, refer to the ESS documentation, as described in “Related information: For more information on the ESS, see these references:” on page 172.



Figure 2-1 Enterprise Storage Server

2.1 Introducing the ESS Web Interface

The Java based Web interface consists of two consoles, the IBM TotalStorage Enterprise Storage Server Specialist (ESS Specialist) and the IBM TotalStorage Enterprise Storage Server Web Copy Services (ESS Web Copy Services).

The ESS Specialist is used for querying the status of the ESS and for displaying, defining and modifying the operation, configuration and administration of the ESS.

The ESS Web Copy Services is used for querying the status of remote operations and performing tasks with the advanced copy functions FlashCopy and Peer to Peer Remote Copy (PPRC) of more than one ESS.

We will cover configurations performed using the ESS specialist.

Accessing the ESS specialist

To access the ESS specialist, we use the Netscape Navigator and point it to the hostname or IP address of any of the two ESS clusters in the Uniform Resource Locator (URL) field. The Web browser then connects to the Web server in the ESS cluster, which handles the requests from the browser.

Since the programs, that run in the Web browser window, are Java applets, the Java Virtual Machine (JVM) used by the browser needs to be capable of supporting the applets. Browsers, that are supported by the ESS Web Interface are:

- ▶ Netscape Navigator
 - 4.04 with JDK 1.1 fixpack
 - 4.05 with JDK 1.1 fixpack
 - 4.06
 - 4.5x
 - 4.7x
- ▶ Microsoft Internet Explorer
 - 4.x with JVM 4.0 or 5.0
 - 5.x with JVM 4.0 or 5.0

The ESS Specialist does not support Netscape Navigator 6.0. Other Browsers might work, but are not currently supported by IBM.

We show a short overview of the configuration panels of the ESS Specialist. These are the panels that are accessed through the buttons displayed on the left of the screen.

2.1.1 Logging on to the ESS Specialist

Here we will detail the steps to log on to the ESS Specialist.

Viewing the Introduction panel

Once connected to one of the ESS clusters, either reloading the Web page, using the **Reload** button of the Web browser, or clicking the **Introduction** button will take us to the introduction panel of the ESS Specialist, as shown in Figure 2-2.

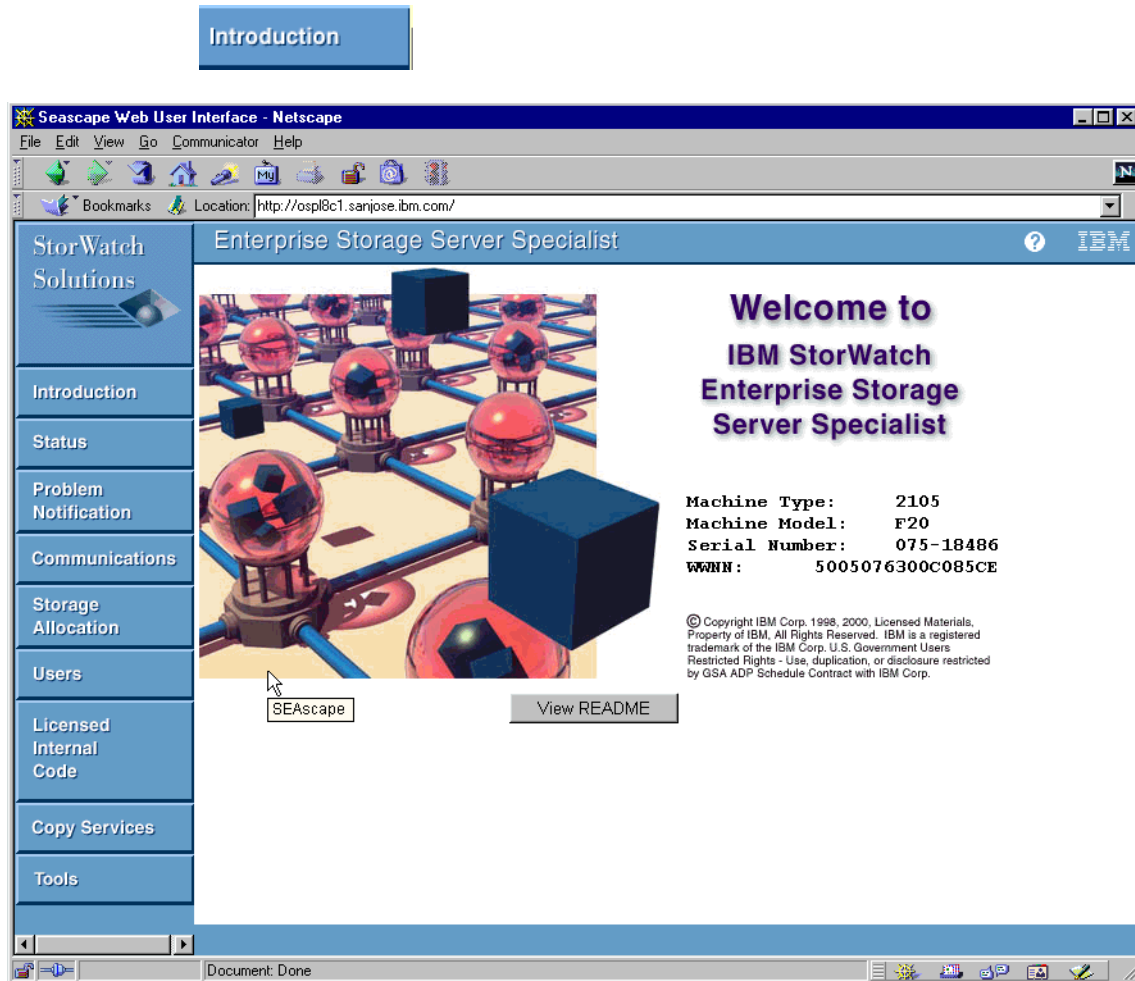


Figure 2-2 Introduction panel of the ESS specialist

To the left, we see the main button panel, from where we can access the viewing and configuration panels of the ESS Specialist. Also, from there we can connect to and troubleshoot the ESS Web Copy Service. On the main window we can see the machine type, machine model, and serial number displayed. Also, we see the identification of the Fibre Channel ESS in the Fibre Channel fabric in the form of its World Wide Node Name (WWNN).

Accepting the ESS site certificate

By clicking on any one of the buttons to the left, we are forced to login to the site. If this is the first time that our browser is connecting to the ESS, the browser will display a sequence of access security panels. These are used to tell the browser that the site that it is pointing to, should be considered as a secure site and that encrypted communication should be allowed.

The first two panels containing information on the site are shown in Figure 2-3 and Figure 2-4.



Figure 2-3 New Site Certificate, start window



Figure 2-4 New Site Certificate, information window

More information on the certificate, including the certificate fingerprint and the time frame for which this certificate is valid, can be displayed by pressing the **More Info** button, as shown in Figure 2-5.

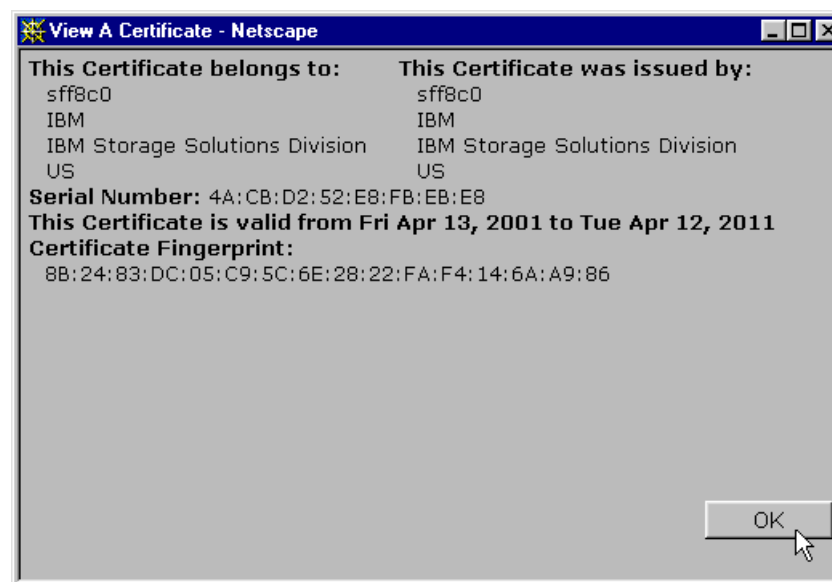


Figure 2-5 New Site Certificate, More Info window

After closing the Information window, we are asked about our decision to accept the certificate or not, and if we would like to accept the certificate until it expires. This is shown in Figure 2-6.

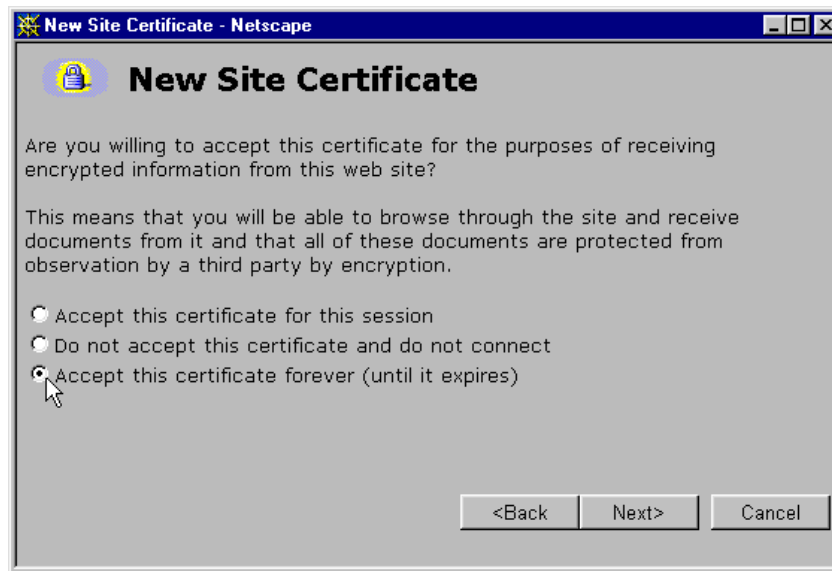


Figure 2-6 New Site Certificate, Accepting the certificate

Logging in with the default user

After completing the certification check, the Login Window, shown in Figure 2-7, will display. The default user is `storwatch` and the default password is `specialist`. Both are case sensitive and must be entered in lower case. For security reasons, it makes sense to change the default user and password to prevent unauthorized access.



Figure 2-7 Site login window

The session will be valid as long as any browser window is open. Therefore, when the work has been performed and security is important, we suggest closing the browser to log off from the ESS specialist.

Be aware that a problem within the browser could keep the browser process running, even when all browser windows are closed. This could lead to the possibility of unauthorized personnel accessing the ESS with the Web interface. A reboot or a new logon to the workstation would solve the problem, but there is another method. This method involves ending the process with the Task Manager. There will be a process **netscape.exe** in the Processes panel, which can be closed with the **End Process** button.

2.1.2 Operating the ESS

Here we introduce the panels that we use to view the status of the ESS, for instance, viewing the problem log and performing operation functions — for example, changing the remote service access and resetting the PE password.

Viewing the Status — Graphical View panel

Using the **Status** button, we get to the Status — Graphical View panel, as shown in Figure 2-8.

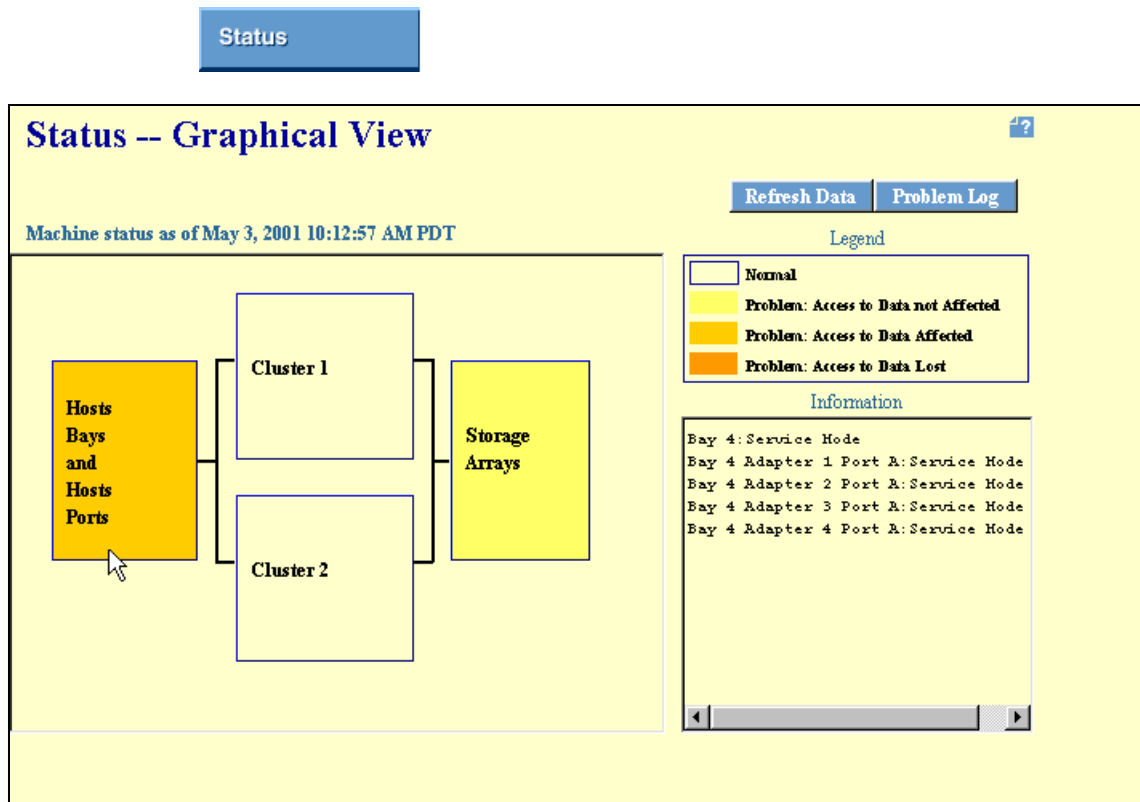


Figure 2-8 Status — Graphical View panel

This panel allows us to quickly check for any errors in the ESS. To demonstrate the features of the Status — Graphical View panel and the Status — Problem Log panel we show them on a different ESS that is currently in service. What we see is a problem detected in the storage arrays, with access to data not affected. The other problem detected within the host bays, however, does affect access to data, which is identified by a darker color.

Viewing the Status — Problem Log panel

To get a detailed view on the problems encountered, we use the **Problem Log** button on the upper right of the screen and get the Status — Problem Log panel, as shown in Figure 2-9.



Figure 2-9 Status — Problem Log panel

Using the **Graphical View** button, we are able to go back to the graphical view.

Viewing the Problem Notification panel

Using the **Problem Notification** button, we get the Problem Notification panel, shown in Figure 2-10.

Problem Notification

Problem Notification

Email and Pager Notification

Email Destinations	Informational	Errors
test@us.ibm.com	no	no

Pager User Name	Pager Number	Pin Number	Enabled	Pager Type
test	9,1-201-440-8777	61398	no	Alphanumeric

SNMP Notification

MIB Read Access - Disabled

SNMP Trap Destination:

Modify Email

Modify Pager

Modify SNMP

Figure 2-10 Problem Notification panel

From within this panel we can configure the various types of problem notification provided with the ESS. The options available are e-mail, pager and Simple Network Management Protocol (SNMP).

Chapter 2. IBM TotalStorage Enterprise Storage Server configuration 85

Viewing the Communications panel

Using the **Communications** button, we get the Communications panel, as shown in Figure 2-11.

Communications

Communications

?

TCP/IP

	Cluster 1	Cluster 2
Hostname	ospl8c0	ospl8c1
IP Address	9.113.20.70	9.113.20.71
Subnet Mask	255.255.254.0	255.255.254.0
Default Router	9.113.21.254	9.113.21.254
Domain	sanjose.ibm.com	sanjose.ibm.com
Domain Name Server	9.113.42.250	9.113.42.250

Remote Support

Call Home
Status - Disabled

Remote Service Access
Status - Disabled

Modify Remote SupportReset PE Password

Figure 2-11 Communications panel

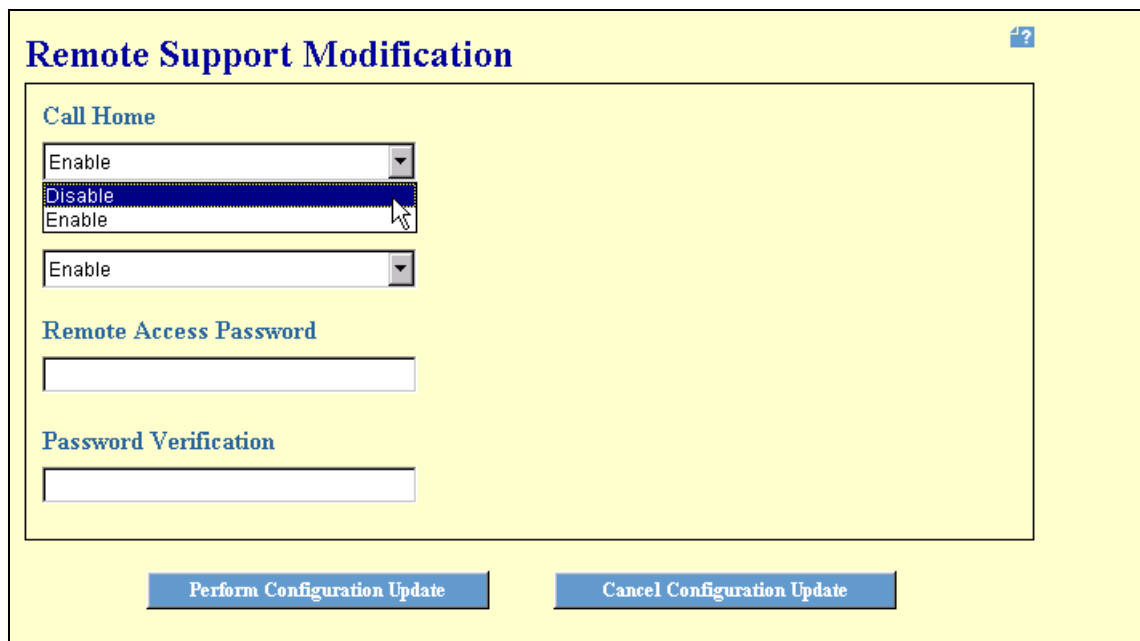
Here we can view the network information for which the ESS is configured, and verify if the Call Home feature and Remote Service Access is enabled. Also, using the **Reset PE Password** button, we can reset the PE password. Only users with administration or configuration levels of authority can modify these panels. The network settings cannot be changed from within the ESS Specialist. An IBM Service Support Representative (SSR) must perform this task using the service terminal.

Generating a new PE password

Resetting the PE password means actually that you generate a password for Product Engineers (PE), which they can use to access the ESS from remote for repair actions. By pressing the **Reset PE Password** button, a warning window will appear, and after clicking the **Yes** button, the window with the new PE password will be displayed. This password grants the PE unlimited access to the ESS LIC and the ESS configuration files for 168 hours (7 days).

Using the Remote Support Modification panel

Using the **Modify Remote Support** button, we get the Remote Support Modification panel, shown in Figure 2-12.



The screenshot shows the 'Remote Support Modification' panel. It has a title bar with a question mark icon. The panel contains three sections: 'Call Home' with two dropdown menus (the first is set to 'Enable' and the second is set to 'Disable'), 'Remote Access Password' with a text input field, and 'Password Verification' with a text input field. At the bottom, there are two buttons: 'Perform Configuration Update' and 'Cancel Configuration Update'.

Figure 2-12 Remote Support Modification panel

The Remote Support Modification panel allows us to enable or disable the Call Home feature and the Remote Service Access. In the event of failure of an ESS component, the Call Home feature enables the ESS to dial to the IBM service center. The Remote Service Access enables an SSR to dial into the ESS to perform a limited set of non-disruptive diagnostic activities.

Viewing the Licensed Internal Code panel

Using the **Licensed Internal Code** button we get the “Licensed Internal Code” panel, as shown in Figure 2-13.

Licensed Internal Code

Licensed Internal Code

View Readme ?

Licensed Internal Code (LIC) Levels

Cluster	LIC Source	Version	Activation Date
Cluster 1	Active LIC	1.4.0.49	2001/05/02
Cluster 1	Previous LIC	1.4.0.44	2001/05/02
Cluster 1	Next LIC	1.4.0.49	
Cluster 2	Active LIC	1.4.0.49	2001/05/02
Cluster 2	Previous LIC	1.4.0.44	2001/05/02
Cluster 2	Next LIC	1.4.0.49	

Licensed Feature Codes

Cluster	Description	Feature Code Capacity Limit	Capacity Used
Cluster 1	Parallel Access Volumes (PAV)	Up to 16.0 TB (S/390)	0.9469 TB (S/390)
Cluster 1	Extended Remote Copy (XRC)	Up to 16.0 TB (S/390)	0.9469 TB (S/390)
Cluster 1	Peer-to-Peer Remote Copy (PPRC)	Up to 16.0 TB (ESS)	2.2095 TB (ESS)
Cluster 1	FlashCopy	Up to 16.0 TB (ESS)	2.2095 TB (ESS)
Cluster 2	Parallel Access Volumes (PAV)	Up to 16.0 TB (S/390)	0.9469 TB (S/390)
Cluster 2	Extended Remote Copy (XRC)	Up to 16.0 TB (S/390)	0.9469 TB (S/390)
Cluster 2	Peer-to-Peer Remote Copy (PPRC)	Up to 16.0 TB (ESS)	2.2095 TB (ESS)
Cluster 2	FlashCopy	Up to 16.0 TB (ESS)	2.2095 TB (ESS)

Figure 2-13 Licensed Internal Code panel

Here we can see the LIC levels of the ESS clusters and the licensed feature codes that include all additional ESS functionality like copy services.

2.1.3 Configuring the ESS

Following, we introduce the panels that serve as entry points to the logical configuration of the ESS and link to the ESS Web Copy Services.

Viewing the Storage Allocation — Graphical View panel

Using the **Storage Allocation** button we get the Storage Allocation — Graphical View panel, shown in Figure 2-14.

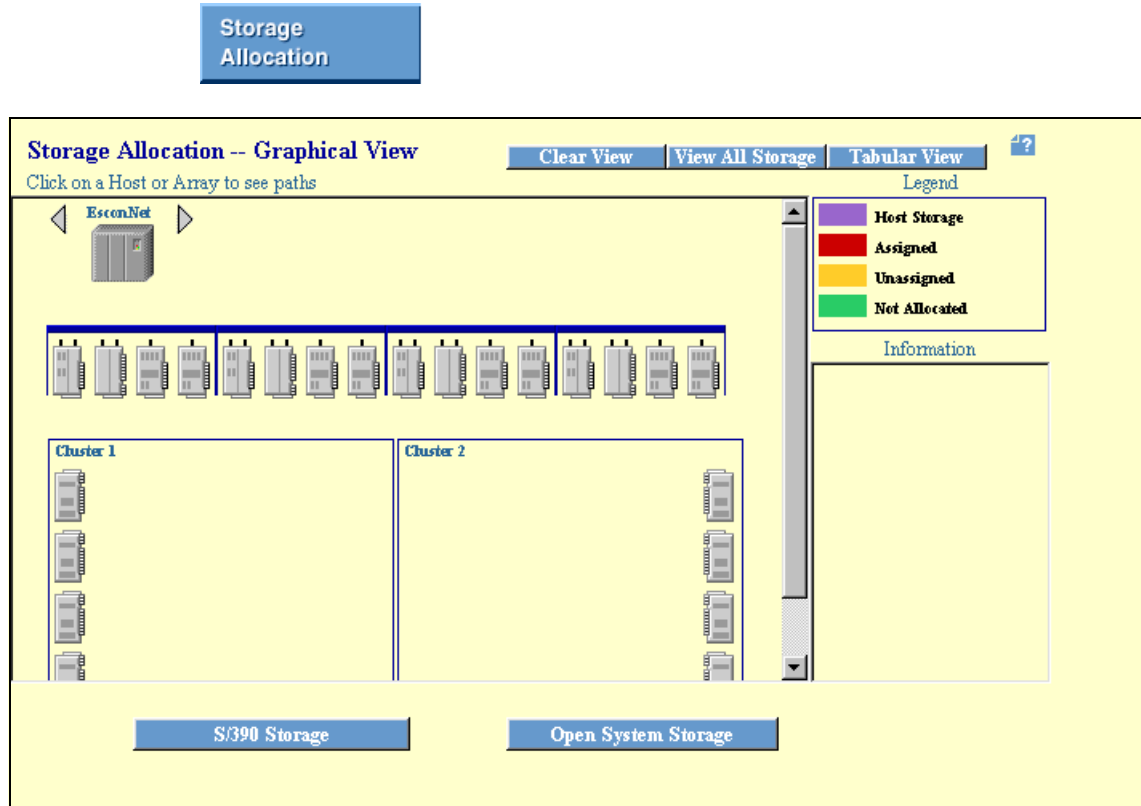


Figure 2-14 Storage Allocation — Graphical View panel

This is the start panel for host and storage based configurations. This panel represents an ESS without anything configured.

Beginning with 2.2, “Starting the Open Systems Storage Configuration” on page 96, we will start with this unconfigured ESS and show how to configure the ESS for our SAN in detail.

Using the Copy Services link

Using the **Copy Services** button we link to the copy services server.



We have to login to gain access, just like we did for the ESS Specialist. This opens the Java based Web interface for the ESS Web Copy Services server in a dedicated Navigator window, connects to it and loads the Java applets, as shown in Figure 2-15.




Figure 2-15 Loading copy services server applets

Once connected, we get the introduction panel for the copy services configuration options, as shown in Figure 2-16.

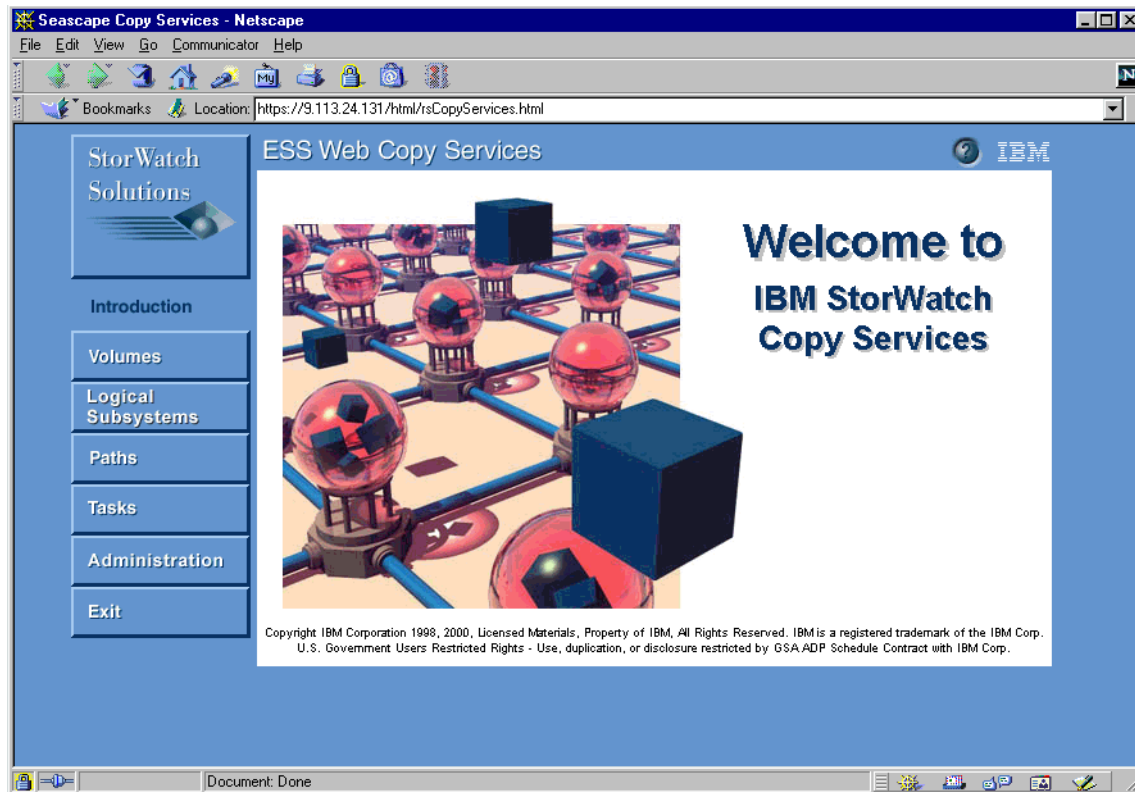


Figure 2-16 ESS Web Copy Services introduction panel

Using the Tools help page

Using the **Tools** button we get a help page with embedded action links for troubleshooting the ESS Web Copy Services.



As mentioned before, the ESS Web Copy Services is a separate server program running on an ESS cluster, which is accessed through a different browser window than the ESS Specialist. However, the ESS Specialist provides this panel for troubleshooting, which makes sense because the ESS Specialist might not be affected from the problem the ESS Web Copy Services server might have.

The helpfile with embedded action links we get is shown in Figure 2-17 and continued in Figure 2-18.

Tools: Copy Services Trouble Shooting and Disaster Recovery

- [Recovering from an 'Unable to Connect to the Copy Services Server' message](#)
- [Determining which Copy Services server is the active server?](#)
- [Reset to primary](#)
- [Reset to backup](#)
- [Disable](#)

Recovering from an 'Unable to Connect to the Copy Services Server' message

An 'Unable to Connect to Copy Services Server' message occurs when the primary Copy Services server does not respond to new activity or when it does not have complete information from all Copy Services clients. When you receive this message take these corrective actions:

1. Sign off from ESS Specialist and close the browser window. Open the browser and sign on to ESS Specialist.
2. Click **Copy Services** from the ESS Specialist Welcome panel.
3. If steps 1 and 2 are unsuccessful, shut down and restart the workstation that is running your browser.
4. Open the browser and access the ESS Specialist again.
5. Click **Copy Services** from the ESS Specialist Welcome panel.

Note: If steps 3 - 5 do not resolve the problem, you need to reset ESS Web Copy Services.

Attention: We recommend that you not perform a reset, unless you are certain that there are no other recovery options. When you perform this function, please be aware that:

- You will lose any PPRC or FlashCopy tasks for which you have not received a successful completion message.
- Established PPRC and FlashCopy relationships are maintained.
- You cannot submit any additional command-line interface (CLI) tasks until ESS Web Copy Services has initialized.

[Reset ESS Web Copy Services](#)

[Top](#)

Determining which Copy Services server is the active server?

The active Copy Services Server is the ESS cluster that is currently running the Copy Services server code.

[Find the active Copy Services server](#)

[Top](#)

Reset to primary

Use the **reset to primary** button to restart Copy Services with the primary server as the active server.

Attention: We recommend that you not perform a reset, unless you are certain that there are no other recovery options. When you perform this function, please be aware that:

- You will lose any PPRC or FlashCopy tasks for which you have not received a successful completion message.
- Established PPRC and FlashCopy relationships are maintained.
- You cannot submit any additional command-line interface (CLI) tasks until ESS Web Copy Services has initialized.
- You lose all the tasks that were created and saved during the time the active server was the backup server.

[Reset ESS Web Copy Services](#)

Reset to backup

Use the **reset to backup** button to restart Copy Services with the backup server as the active server.

Figure 2-17 Tools help page, Web Copy Services trouble shooting

Attention: We recommend that you not perform a reset, unless you are certain that there are no other recovery options. When you perform this function, please be aware that:

- You will lose any PPRC or FlashCopy tasks for which you have not received a successful completion message.
- Established PPRC and FlashCopy relationships are maintained.
- You cannot submit any additional command-line interface (CLI) tasks until ESS Web Copy Services has initialized.

[Reset ESS Web Copy Services](#)

Disable

Use the **disable** button to disable Copy Services.

Attention: We recommend that you not disable Copy Services, unless you are certain that there are no other options. When you perform this function, please be aware that:

- You will lose any PPRC or FlashCopy tasks for which you have not received a successful completion message.
- Established PPRC and FlashCopy relationships are maintained.
- You cannot submit any additional command-line interface (CLI) tasks until ESS Web Copy Services has initialized.

[Disable ESS Web Copy Services server](#)

[Top](#)

For information on Trouble Shooting and Disaster Recovery, see [Copy Services Trouble Shooting and Disaster Recovery](#) .

Figure 2-18 Tools help page, Web Copy Services trouble shooting, continued

2.1.4 Administer the ESS

We will now introduce the administration of users on the ESS.

Using the User Administration panel

Using the **Users** button we get the User Administration panel, as shown in Figure 2-19.



User Administration

User Permissions

Perform Sort

Total Users: 1

no sort

no sort

no sort

no sort

Username	Access Level	IP Address Range	User Comment
storwatch	Administration		DEFAULT ADMIN ID, WILL BE REMOVED

Modify Users

Figure 2-19 User Administration panel

In our case there are no users configured, other than the default user that we logged on with.

Using the Modify Users panel

In the Modify Users panel, as shown in Figure 2-20, we can add users with different access levels.

Modify Users

User Account

Name

Password

Password Verification

Access Level

Administration

IP Address Range (Optional)

Comments (Optional)

Add >>

<< Remove

User List

Username	Access Level	IP Address Range
storwatch	Administration	

Perform Configuration Update

Cancel Configuration Update

Figure 2-20 Modify Users panel

The access levels are:

View	Viewing the status and the configuration.
Operation	Viewing the status and the configuration and performing operation functions.
Configuration	Viewing the status and the configuration, performing operation and configuration functions.
Administration	Viewing the status and the configuration, performing operation, configuration and administration functions.

After the first administrative user is defined on the ESS, the default user will be removed. If all administration users have been, or get, deleted, the default user will be restored after 30 minutes.

2.2 Starting the Open Systems Storage Configuration

In the topics that follow, we will guide users through the storage and Fibre Channel environment related parts as we progress through and build the configuration. The tasks performed here will affect the way the ESS presents itself in our SAN and to the SAN attached Fibre Channel hosts.

2.2.1 Viewing the storage configuration status

The panels described here will always show the status of the configuration of the ESS. There are two panels, one that illustrates the configuration graphically, and one that shows it in an tabular view.

Using the Storage Allocation — Graphical View panel

Using the **Storage Allocation** button we get the Storage Allocation — Graphical View panel, as shown in Figure 2-21.

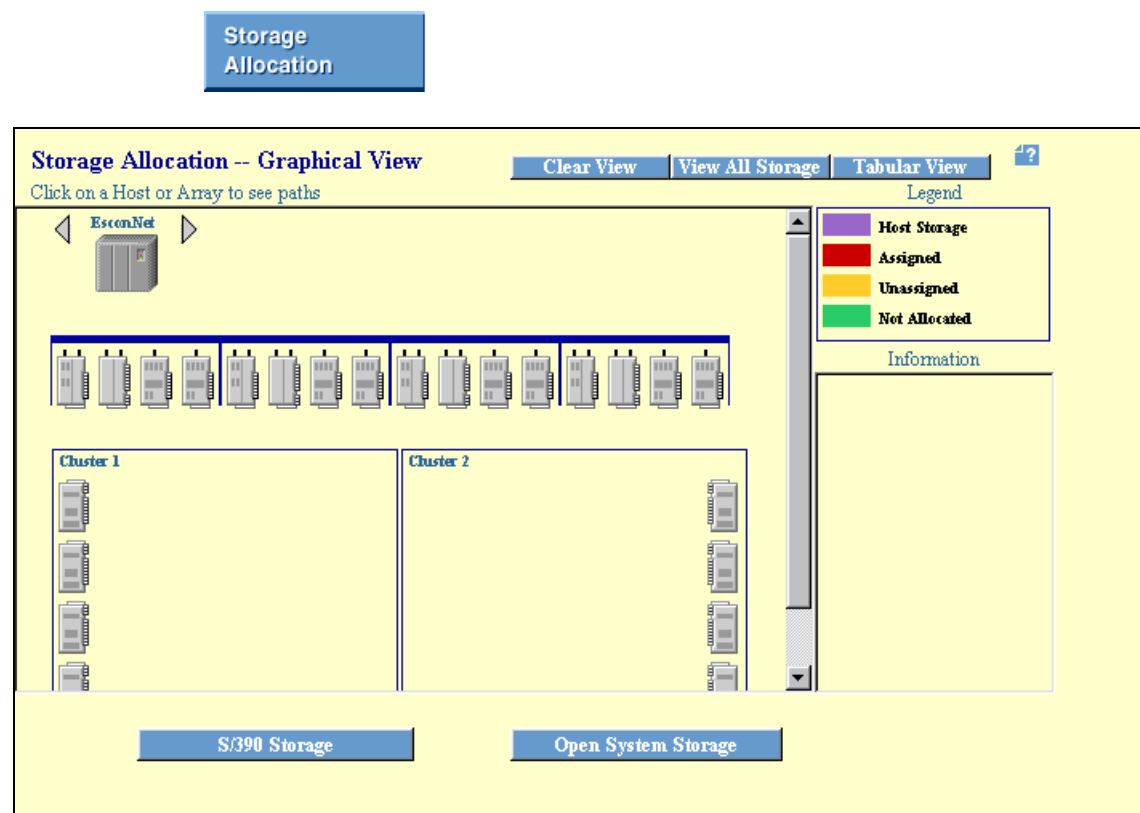


Figure 2-21 Graphical View, unconfigured ESS with ESCON, SCSI and FC

In this panel, we get the logical view and the status of the installed and configured storage related components of the ESS. These are the parts of the ESS that need to be configured for ESS logical volumes to be accessible from host FC ports.

The panel is interactive, which means, as we select different entities on this panel, the panel will reflect the configuration of the ESS according to the selected entity. Those can be the host interface icons, the ESS interface icons, the SSA device adapter icons and the disk group rectangles. Also, the **View All Storage** button allows us to see the status of all configured disk groups at once

View All Storage

To start with our ESS is unconfigured, and in the screen shots that follow we will show how the panel will change.

Using the host interfaces row

The top row, shown in Figure 2-22, presents symbols for the host interfaces, that have access to the ESS.



Figure 2-22 Host interfaces row, without any configured host interface

In the case of hosts with Fibre Channel connectivity, there will be an icon like the one shown in Figure 2-23, for each host FC port that is defined in the ESS with its World Wide Port Name (WWPN). The same applies to defined SCSI attached hosts.



Figure 2-23 Open systems host port icon

There is no host FC port configured at this time, therefore, in Figure 2-22, there is no icon shown. For ESCON attachment, if installed, there is a placeholder icon for all connections. The host interface icons can be selected, which will change the view of the panel to represent the configuration according to the selected host interface.

Using the ESS interfaces row

Below the host interface row, we see the row with icons for the ESS installed interface adapters, shown in Figure 2-24.




Figure 2-24 Installed ESS interfaces row

The row is divided into four columns to show the location in the ESS interface bays. Our ESS is equipped with four ESCON adapters and four SCSI adapters, with one of them in each bay, and 8 Fibre Channel adapters, with two of them in each bay. For a Fibre Channel ESS, which also has to serve a few SCSI and ESCON attached hosts, this is a fairly balanced setup.

As opposed to the host interfaces row, the icons shown in the ESS interfaces row represent adapters, not ports. If an adapter type, for instance a SCSI adapter, features two ports, we have still one icon.

We can also see how many ports a specific adapter type offers. We can see that ESCON and SCSI adapters are dual ported and the FC adapters are single ported. For clarity, we use the term ESS FC port in this redbook, even when talking about the adapter itself.

The ESS interface adapter icons can also be selected, which will change the view of the panel. For instance, it will display information of the type and configuration of the adapter.

Graphically, the representation of the adapters is different to help to distinguish between them. The different representations of the adapters are shown in Figure 2-25, Figure 2-26, and Figure 2-27.




Figure 2-25 ESS ESCON adapter with two ports




Figure 2-26 ESS SCSI adapter with two ports




Figure 2-27 ESS FC adapter with one port

Viewing the association between host ports and ESS ports

In the ESS, we do not tie a host FC port to a Fibre Channel port in the ESS, which was what we did with the SCSI adapters. Every host FC port within the Fibre Channel fabric will be identified by its WWPN and can access data through every ESS FC port unless other measures are taken. We will see later during the ESS configuration, that we can restrict host FC ports through which ESS FC port it will be able to access ESS logical volumes.

To show the independence of the port, the selected host FC port is graphically connected to each ESS FC port, as shown in Figure 2-28. In this example, we will assume that we already have two host FC ports configured to illustrate this point.




Figure 2-28 FC host port selected

In contrast to this, the host SCSI ports have an affinity to the ESS SCSI ports, which is visible by selecting a host SCSI port, as shown in Figure 2-29.




Figure 2-29 SCSI host port selected

Using the disk group section

The two ESS clusters with the SSA device adapters and the SSA loops with the disk groups are shown below the ESS Fibre Channel ports.

This is shown in Figure 2-30 and is how it looks if there is no disk group configured.

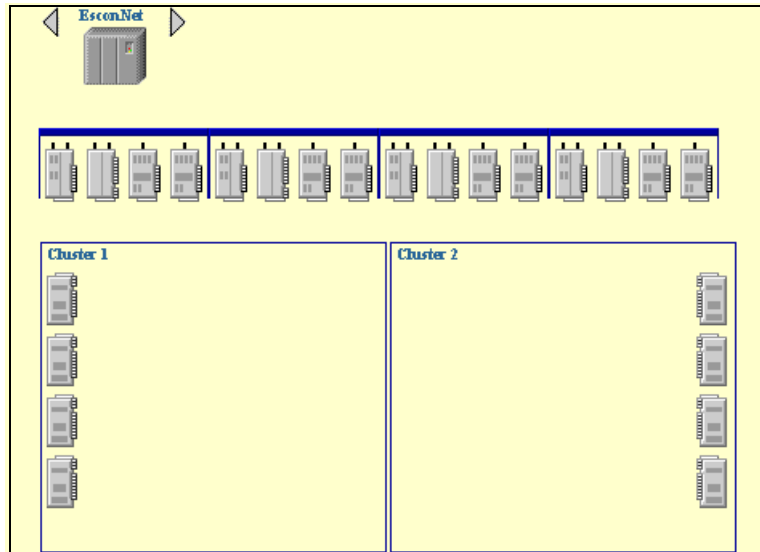


Figure 2-30 SSA device adapters without any disk group configured

Again, we are anticipating the configurations that are to be made later. To graphically show the state the disk groups are in, we press the **View All Storage** button on the upper right of the panel.

This is shown in Figure 2-31. Disk groups can also be selected to show the state of the selected storage. Then, the host interface row will only show the host FC ports, that have ESS logical volumes assigned to that exist on this particular disk group.

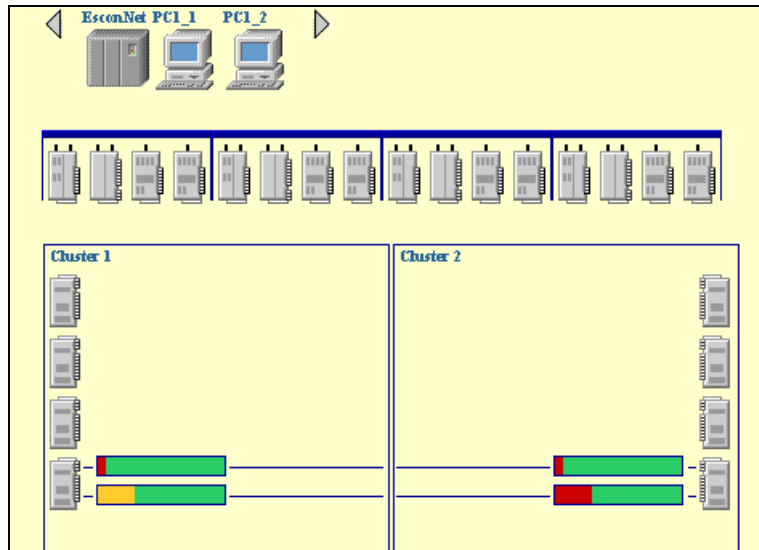


Figure 2-31 Disk groups, view all storage mode

To show the state of the disk groups, which contain ESS logical volumes assigned to a particular host FC port, we select the host FC port first and then one of the yellow highlighted ESS FC ports, which is shown in Figure 2-32.

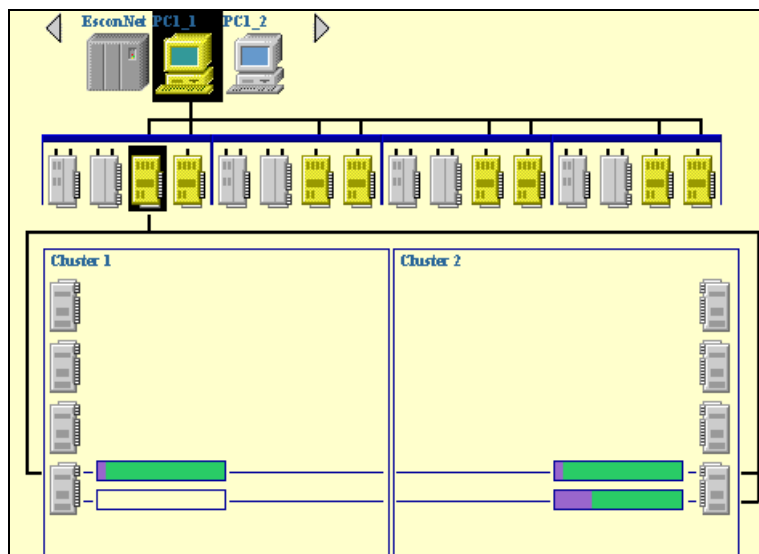


Figure 2-32 Disk groups, storage assigned to selected host FC port

There are different colors to indicate the state of the disk groups. At the upper right of the screen, shown in Figure 2-33, a legend is displayed, which explains what the different colors in the disk groups represent.




Figure 2-33 Legend of the disk group status colors

The colors indicate the following:

- Purple, Host Storage** Disk storage that contains ESS logical volumes that are assigned to a specific (selected) host FC port. This displays only if a host FC port is selected first.
- Red, Assigned** Disk storage that contains ESS logical volumes that are assigned to other than the selected host FC port, or that are assigned to any host FC port, when no host FC port was selected. This happens, for example, when we select a specific disk group or when we press the **View All Storage** button.
- Yellow, Unassigned** Disk storage that contains ESS logical volumes that were assigned to host FC ports before and then were unassigned. Those ESS logical volumes can be assigned to any host FC ports.
- Green, Not Allocated** Disk storage that contains no ESS logical volumes. This storage can be used to define and assign ESS logical volumes.

We will explain more details of the Storage Allocation — Graphical View panel later on as we progress with our configuration and the panel shows the information relating to our configuration changes.

Using the Storage Allocation — Tabular View panel

To get a detailed view of the ESS logical volumes later on, we will use the **Tabular View** button to get to the view as shown in Figure 2-34.



Storage Allocation -- Tabular View

List of Assigned Volumes

Print Table

Perform Sort

Graphical View

no sort

no sort

no sort

no sort

no sort

no sort

no sort

no sort

Host/SSID	LSS/LCU	Volume	Type	Size	Host Adapter	Location	Shared

Figure 2-34 Tabular View, no ESS logical volumes configured

As expected, the table that contains the detailed description of all volume assignments made, is empty right now.

2.2.2 Using the Open System Storage panel

To start with the configuration of the ESS, we press the **Open System Storage** button, as shown in Figure 2-35. This is the door to the configuration of open systems storage. It will lead us to the Open System Storage panel, as shown in Figure 2-36. The **S/390 Storage** button is also accessible, because there are ESCON cards installed.

S/390 Storage

Open System Storage

Figure 2-35 Entry for open systems storage configurations

Using the Open System Storage Panel

The Open System Storage panel consists of the Hosts Systems table, the Assigned Volumes table and the configuration panel buttons. We will see shortly that those two tables are showing the defined host FC ports and, to the selected host FC port, the assigned ESS logical volumes, as shown in Figure 2-36.

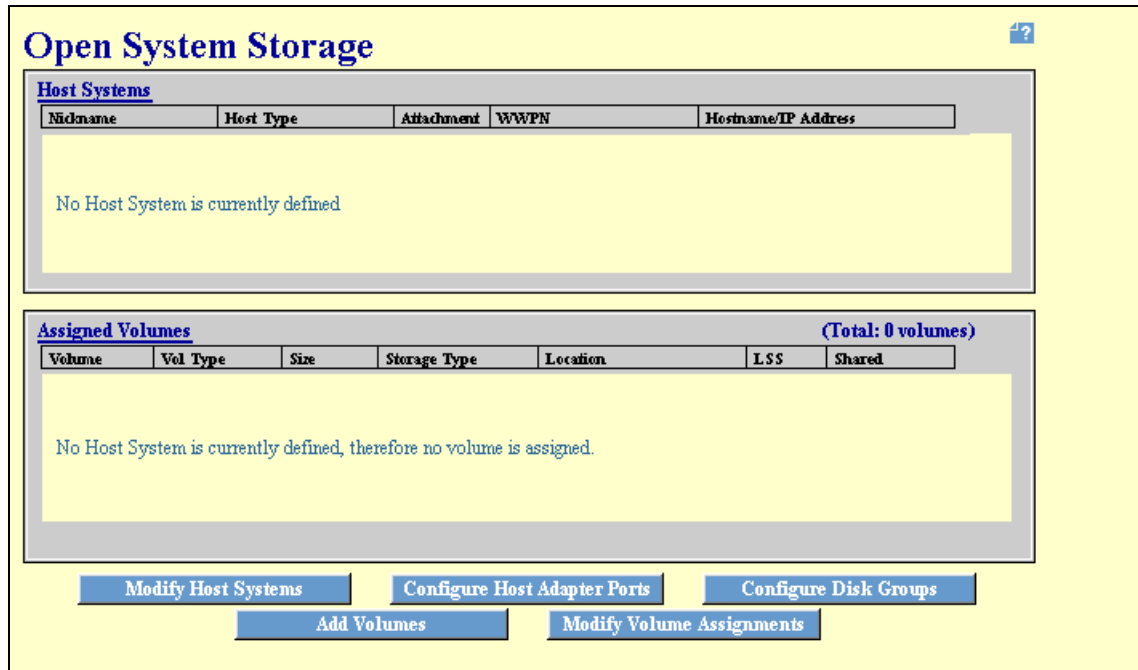


Figure 2-36 Open System Storage panel, no host FC ports defined

The configuration panel buttons at the bottom are the entry points to all of our storage related configuration tasks. The Open System Storage panel is like a hub, from where we start every open systems storage configuration. In this section, we will indicate our use of the configuration panel buttons by pointing the mouse on the button.

Using the proper order of configuration steps

This is an introduction to the dependencies that some configuration steps have. Our ESS is not configured, so we show where we start the configuration. What we want to achieve is to prepare the ESS to provide some ESS logical volumes to host FC ports. Some of the configuration steps have to be done prior to others.

Figure 2-37 shows the dependencies of the configuration steps.

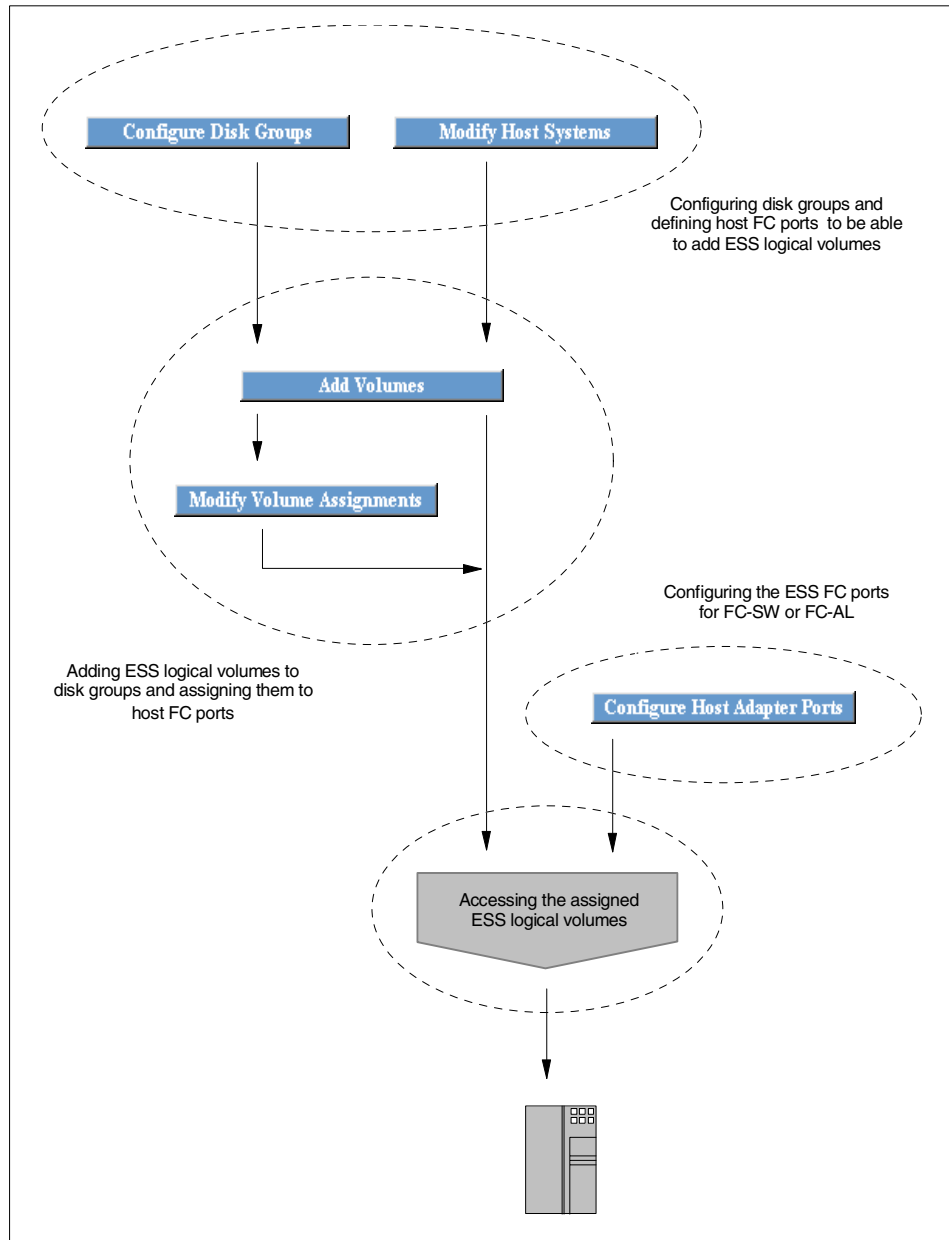


Figure 2-37 ESS configuration dependencies

One step that is independent from the others, is the configuration of the ESS FC ports for FC-SW or FC-AL. For this we have to know what type of FC topology we want to connect with a particular ESS FC port. Other steps rely on the completion of others.

In order to add ESS logical volumes to disk groups, we need to perform two configuration steps. Those are the configuration of some disk groups for RAID or JBOD and the definition of one or more host FC ports by its host type and WWPN. The host FC port definition is needed because the process of adding volumes consists not only of creating the volumes but also of assigning them to a host FC port.

If there are existing ESS logical volumes, we can also assign them to other host FC ports. Now, if both requirements have been met (ESS logical volumes have been assigned and the ESS FC ports are defined), we can access the volumes from the Fibre Channel attached host. Of course, the connectivity must be there and the host system has to be prepared also. How to achieve this is discussed in other parts of this book.

2.3 Configuring disk groups

Now we will show how to configure the ESS for FC attachment in detail. We choose to start our configuration with the disk groups, followed by the configuration of the host FC ports. The completion of both configuration steps enables us to add the ESS logical volumes later on.

2.3.1 Performing the disk group configuration

Pressing the **Configure Disk Groups** button takes us to the Fixed Block Storage panel, as shown in Figure 2-38.

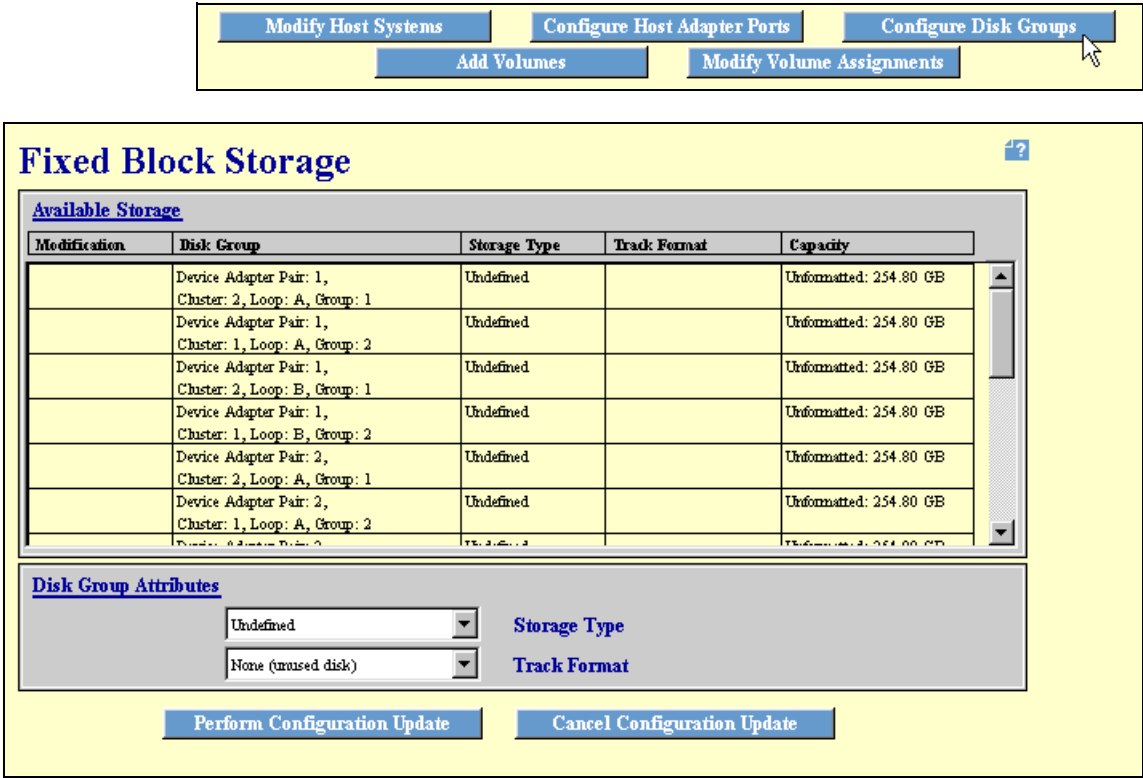


Figure 2-38 Fixed Block Storage panel, no disk groups configured

Using the Fixed Block Storage panel

The Available Storage table lists all of the available physical storage in the ESS with its location, storage type, track format, and capacity. As we can see, when we scroll down the table, no disk array has been formatted at this time.

Note: The disk groups should not be mistaken for the eight packs, which are the physical disk packages, that are installed in the ESS. The disk groups are logical groups that consist of disks from two eight packs, four disks from each of the two. This is the reason why it is only possible to upgrade the ESS in steps of two eight packs.

The unformatted size of the disk groups is 254.8 GB. This is equivalent to seven disks of a capacity of 36.4 GB each. In the ESS, we have two spare disks per SSA loop. In our ESS we have sixteen disk groups, four per device adapter pair and two per SSA loop. This is the reason why every one of our disk groups contains one spare disk. This makes up the total unformatted capacity of 254.8 GB per disk group.

Defining two disk groups for RAID 5

We are going to define the disk groups connected to the A loop of the device adapter pair 4 as RAID array. To do this, we select the first disk group we want to configure and select **Raid Array** from the Storage Type pull-down menu. The storage type, as shown in Figure 2-39, can be RAID Array, or non-RAID. Non-RAID means the disks are going to be configured as “just a bunch of disks” (JBOD).




Figure 2-39 Open systems storage types

Doing this also changes the Track Format to Fixed Block (FB) in the Track Format pull-down menu. After doing this, we select the second disk group in the other cluster, and also select **Raid Array**.

The Modification column of the panel, as shown in Figure 2-40, shows that we have made changes (Defined), and the Capacity column shows that our disk groups are still unformatted. In this case, formatted to the ESS means it is not formatted as either a RAID array or as non-RAID.

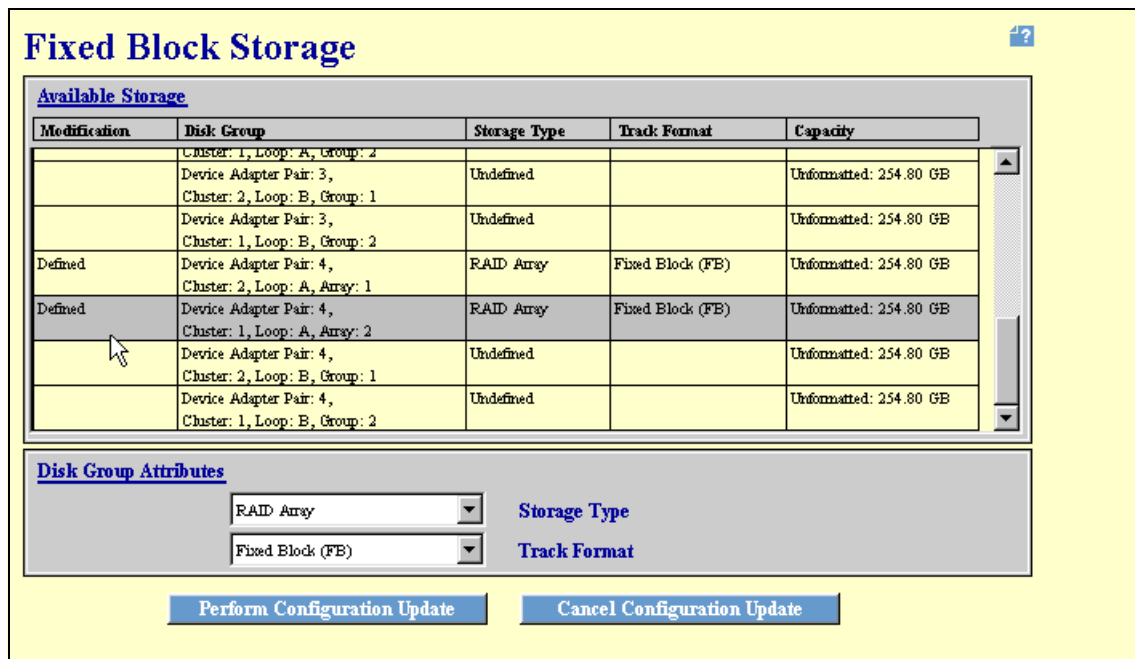


Figure 2-40 Disk groups defined for RAID 5

Performing the configuration update

To apply the changes made, we press the **Perform Configuration Update** button.



A warning message appears, stating that this will be a time consuming action, as shown in Figure 2-41.



Figure 2-41 Warning window, time consuming action

How time consuming this will be depends on how many disk groups are to be configured at once. Pressing the **OK** button executes the script which changes the configuration and gives us a progress window, where we can see which arrays are initializing, as shown in Figure 2-42.




Figure 2-42 Progress window, RAID configuration

After completing the configuration, the ESS Specialist informs us of the success, as shown in Figure 2-43.




Figure 2-43 Status window, successful disk group configuration

We have to press the **OK** button here, which takes us back to the Open System Storage panel.

Viewing changes on the Graphical View panel

From here we can continue with our configuration. To see how the Storage Allocation — Graphical View panel now looks, we press the **Storage Allocation** button on the left of the screen.



The panel shown in Figure 2-44 now illustrates the representation of the two disk groups between the SSA device adapter pair four in loop A. It is the view of the panel with the disk group in cluster one selected.

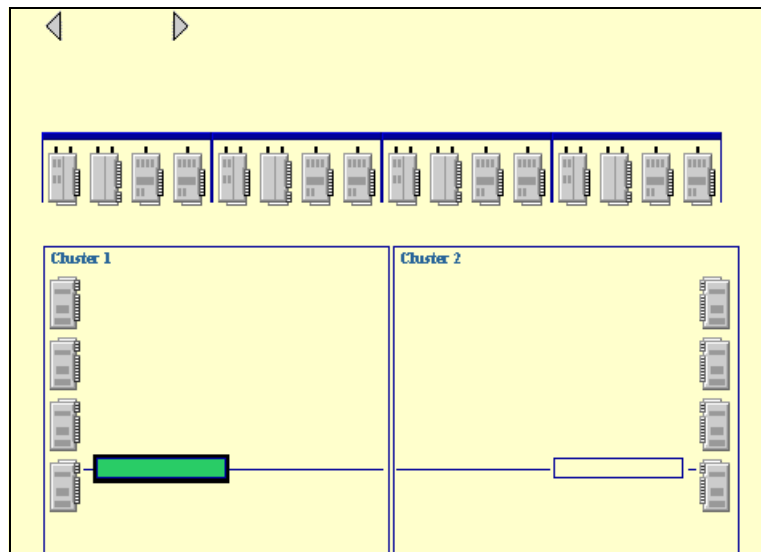


Figure 2-44 Disk group section, configured disk group selected

The rectangle representing the disk group is green to show that this is empty storage. We also see that the host FC port icon row is empty. When selecting a particular disk group, the row will only present and highlight those host FC port icons that have ESS logical volumes assigned to it in the selected disk group. Because there is no host FC port defined, there cannot be any ESS logical volume defined and assigned.

When the disk group is selected, the Information field to the right of the panel shows specific information such as the location of the disk group, the type of storage the disk group is configured for and the available and allocated capacity.

This is shown in Figure 2-45.

Information	
Cluster:	1
Adapter:	4
Loop:	A
Disk Group:	2
Storage Type:	RAID Array
Track Format:	Fixed Block (FB)
Host Storage	0.00/210.48GB (0%)
Assigned	0.00/210.48GB (0%)
Unassigned	0.00/210.48GB (0%)
Not Allocated	210.45/210.48GB (100%)

Figure 2-45 Information window, for RAID 5 configured disk group

Here we see that the capacity of the disk groups that are configured as RAID array, has changed to 210.48 GB. This is mainly because of the space for the redundancy information that is now spread over the disks in the disk group and because of the RAID configuration information.

Defining two more disk groups for RAID 5

We will configure two more disk groups and so we therefore navigate back to the Fixed Block Storage panel.

In Figure 2-46 we show what the Available Storage table looks like with the two disk groups in loop A of the fourth adapter pair for RAID configured. This is prior to performing the configuration update of the next two disk groups.

Fixed Block Storage

Available Storage

Modification	Disk Group	Storage Type	Track Format	Capacity
	Cluster: 1, Loop: A, Group: 2			
	Device Adapter Pair: 3, Cluster: 2, Loop: B, Group: 1	Undefined		Unformatted: 254.80 GB
	Device Adapter Pair: 3, Cluster: 1, Loop: B, Group: 2	Undefined		Unformatted: 254.80 GB
	Device Adapter Pair: 4, Cluster: 2, Loop: A, Array: 1	RAID Array	Fixed Block (FB)	Formatted: 210.48 GB
	Device Adapter Pair: 4, Cluster: 1, Loop: A, Array: 2	RAID Array	Fixed Block (FB)	Formatted: 210.48 GB
Defined	Device Adapter Pair: 4, Cluster: 2, Loop: B, Array: 1	RAID Array	Fixed Block (FB)	Unformatted: 254.80 GB
Defined	Device Adapter Pair: 4, Cluster: 1, Loop: B, Array: 2	RAID Array	Fixed Block (FB)	Unformatted: 254.80 GB

Disk Group Attributes

RAID Array

Storage Type

Fixed Block (FB)

Track Format

Perform Configuration Update

Cancel Configuration Update

Figure 2-46 Two disk groups configured, two defined and two undefined

After we have performed the configuration update, we go back to the Storage Allocation window to see how it now looks as a result of our changes.

Viewing changes on the Graphical View panel

Device adapter pair number four and the connected SSA loops now look like that shown in Figure 2-47.

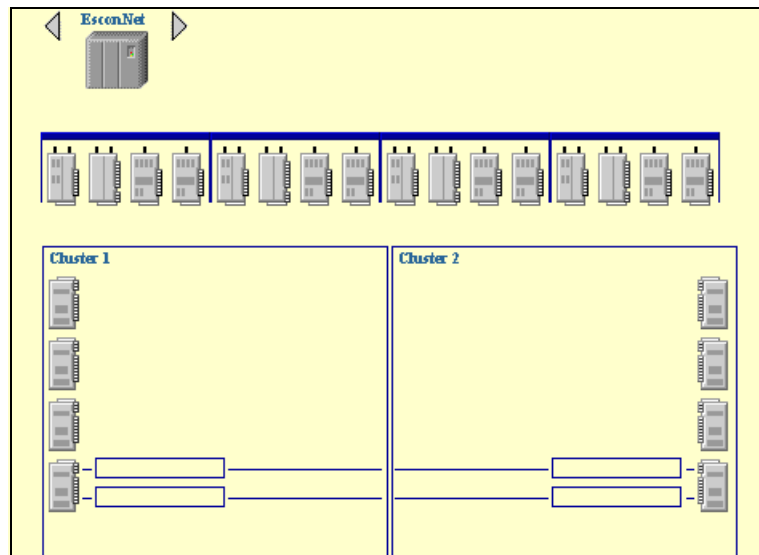


Figure 2-47 Four disk groups configured

Undefining disk groups or redefining disk groups

It may be necessary to reconfigure disk groups. They can be unconfigured or configured to another storage type. If there is a disk group configured for RAID 5, it can be reformatted to JBOD or unconfigured. Doing this is the same process as we did before.

First we select the disk group in the Available Storage table and then we have to choose **Undefined** or **non-RAID** in the Disk Group Attributes field. After we get a warning that all volumes on the selected disk group will be destroyed, we can select the next disk group to be reconfigured. After changing the attributes of the disk groups we perform the update.

A window pops up, informing us that the process will take a while and that we can view the progress on the Status panel. This is with four disk groups in the process of formatting, as illustrated in Figure 2-48.

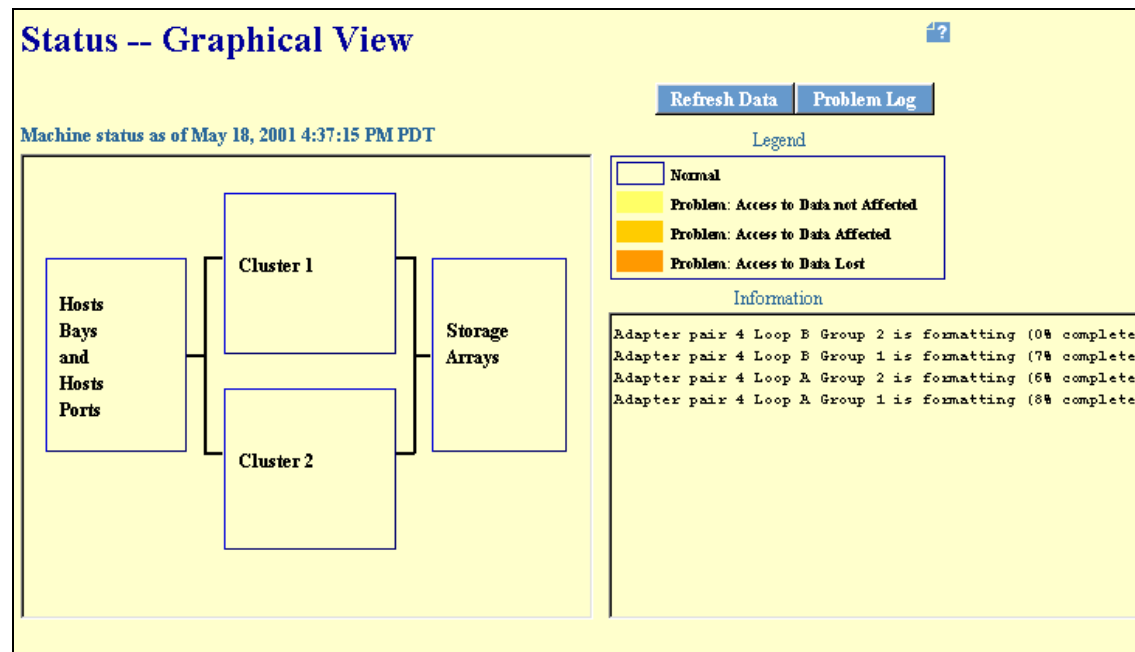


Figure 2-48 Status — Graphical View, four disk groups in the process of formatting

This discussion was just intended to show how it is done. Now we continue from the point where we defined the disk groups.

2.4 Creating and modifying host FC port definitions

To fulfil the second condition to be able to define and assign ESS logical volumes, we will define the host FC ports that are installed in the Fibre Channel attached hosts.

2.4.1 Creating host FC port definitions

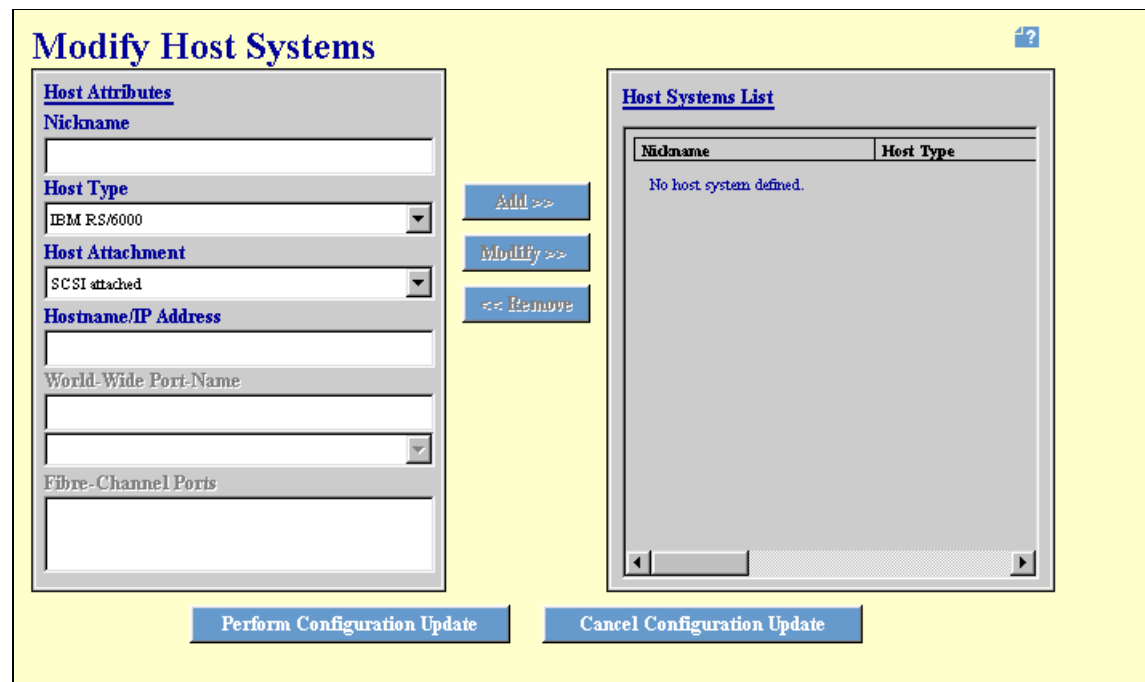
From the Open System Storage panel, we press the **Modify Host Systems** button, as shown in Figure 2-49.



Figure 2-49 Entry to the Modify Hosts Systems panel

Viewing the Modify Host Systems panel

We are now taken to the Modify Host Systems panel, as shown in Figure 2-50. This panel consists of the Host Attributes entry fields and The Host Systems List table. The Host Systems List represents each defined host FC port. If a host has two Fibre Channel ports, this appears as two entries in the Host Systems List. As we can see, there are no host FC ports defined at this time.



Host Attributes	
Nickname	
Host Type IBM RS/6000	
Host Attachment SCSI attached	
Hostname/IP Address	
World-Wide Port-Name	
Fibre-Channel Ports	

Add >>

Modify >>

<< Remove

Host Systems List	
Nickname	Host Type
No host system defined.	

Perform Configuration Update

Cancel Configuration Update

Figure 2-50 Modify Host Systems panel, no host FC ports defined

Creating two host FC port definitions

In the Host Attributes entry fields, we identify the host FC port that will access the ESS. The first entry is a nickname to identify the port in the ESS. Also, we have to choose the host type, which is a PC server; the type of the host attachment, which is Fibre Channel, and the WWPN of the host FC port, as shown in Figure 2-51 and Figure 2-52.

We are doing this for the two host FC ports that we want to access ESS logical volumes, nicknamed PC1_1 and PC1_2.

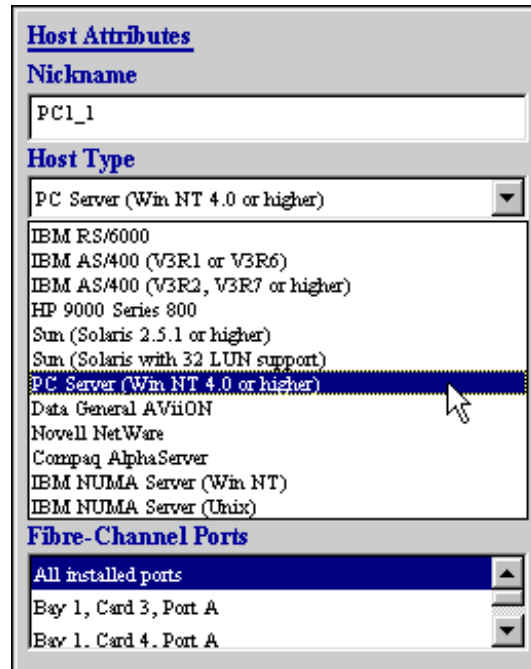


Figure 2-51 Host Attributes frame, host type selection

Host Attributes

Nickname

Host Type

Host Attachment

 (highlighted by mouse)

World-Wide Port-Name

Fibre-Channel Ports

Figure 2-52 Host Attributes frame, host interface type selection

The IP Address is only used by the StorWatch Enterprise Storage Server Expert. It is used to show which host FC ports are in a specific IP host. We will not be using the Expert for our example, so we leave the entry field empty.

By pressing the **Add** button, we add the specified host FC port to the Host Systems List. We do the same for a second host FC port, as shown in Figure 2-53.



Host Systems List					
Nickname	Host Type	Attachment	WWPN	Hostname/IP Address	FC Ports
PC1_1	PC Server (Win NT 4.0 or higher)	Fibre Channel attached	200000E08B00C5B		All installed ports
PC1_2	PC Server (Win NT 4.0 or higher)	Fibre Channel attached	210000E08B017862		All installed ports

Figure 2-53 Host Systems List, host FC ports defined

This host FC port can belong to the same system, or we can define ports for other systems.

Performing the configuration update

Adding ports to the list puts together a script in the background which is executed by pressing the **Perform Configuration Update** button.

Perform Configuration Update

A progress bar informs us about the steps that are being executed, as shown in Figure 2-54.

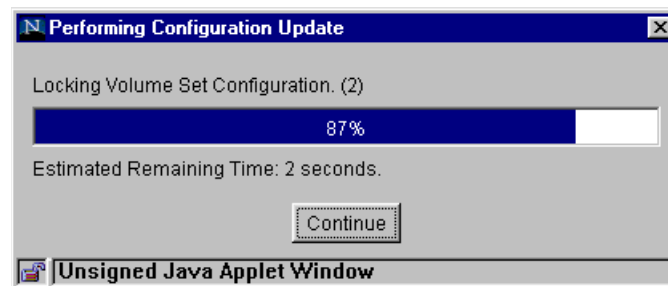


Figure 2-54 Performing the port definition

A window appears which indicates that the configuration update was successful, as shown in Figure 2-55. If the update had failed for any reason, we would get an error message.



Figure 2-55 Successfully added host FC port

Pressing the **OK** button on this window will immediately take us back to the Open System Storage panel.

Viewing changes on the Open System Storage panel

We now see the newly configured host FC ports in the Host Systems Table, as shown in Figure 2-56.

Open System Storage

Host Systems

Nickname	Host Type	Attachment	WWPN	Hostname/IP Address
PC1_1	PC Server (Win NT 4.0 or higher)	FC	200000E08B00CC5B	
PC1_2	PC Server (Win NT 4.0 or higher)	FC	210000E08B017862	

Assigned Volumes

(Total: 0 volumes)

Volume	Vol Type	Size	Storage Type	Location	LSS	Shared
There are no volumes assigned to the currently selected host system						

Modify Host Systems
Configure Host Adapter Ports
Configure Disk Groups

Add Volumes
Modify Volume Assignments

Figure 2-56 Fibre Channel host adapter port; no volume assigned

As we select one of them, we see that there are no ESS logical volumes currently assigned to the host FC ports.

Viewing changes on the Graphical View panel

Now, we would like to see what the Storage Allocation — Graphical View panel looks like. This is shown in Figure 2-57.

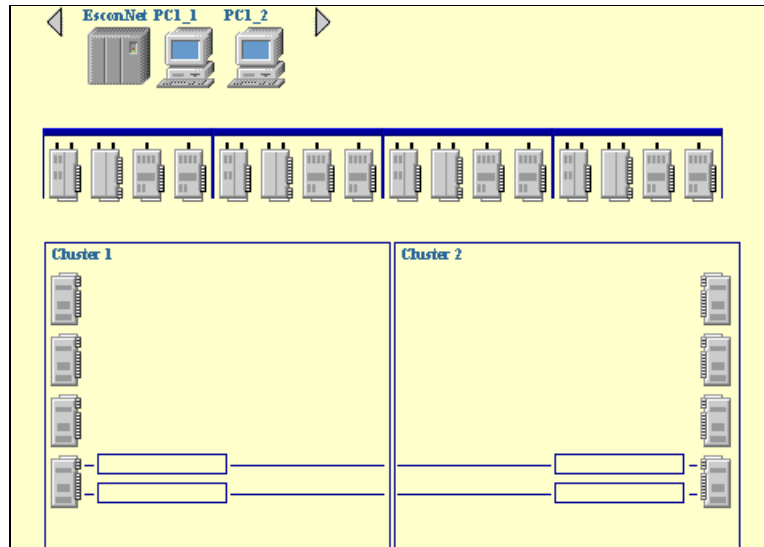


Figure 2-57 Four disk groups, two host FC ports

The two host FC ports we just have configured, now show up with their nicknames in the host interface row. Although we could have used longer nicknames, they would have been truncated, which would make them very hard to identify. The only way would be to click on them and catch their names along with other information like the WWPN and the type of the host system, on the Information window shown in Figure 2-58.

Information	
Nickname:	PC1_1
WWPN:	200000E08B00CC5B
Attachment:	FC
Description:	
Operating System:	PC Server (Win NT 4.0 or higher)

Figure 2-58 Information window, host FC port in PC server

Clicking on the host FC port icon also shows that the port would be able to access volumes through each of the eight installed ESS FC ports if they were all connected to the SAN and no zoning in the fabric would prevent them from accessing the ESS FC ports.

This is illustrated in Figure 2-59, with lines drawn from the selected host FC port to each installed ESS FC port. The involved components also get highlighted yellow.




Figure 2-59 Selected host FC port, no ESS FC port restrictions

2.4.2 Modifying a host FC port definition

Each one of the settings that we applied before, for instance the nickname or the WWPN, can be modified from within the Modify Host Systems panel.

We will now restrict the host FC port PC1_2 to the ESS FC port 3 in bay 1, as an example of modifying the settings for host FC ports from within the Modify Host Systems panel.

Restricting the ESS FC port access of host FC port PC1_2

To do this, we select the host FC port **PC1_2** in the Host Systems List table, as shown in Figure 2-60. This enables us to change the settings for this port in the Host Attributes section.

Modify Host Systems

Host Attributes

Nickname
PC1_2

Host Type
PC Server (Win NT 4.0 or higher)

Host Attachment
Fibre Channel attached

Hostname/IP Address

World-Wide Port-Name
210000E08B017862

Fibre-Channel Ports
All installed ports
Bay 1, Card 3, Port A
Bay 1, Card 4, Port A

Add >>

Modify >>

<< Remove

Host Systems List

Nickname	Host Type
PC1_1	PC Server (Win NT 4.0
PC1_2	PC Server (Win NT 4.0

Perform Configuration Update

Cancel Configuration Update

Figure 2-60 FC port PC1_2 selected for modification

As we show here, the Fibre-Channel Ports list indicates the current setup, which is not restricted in ESS FC ports. It lists all ESS FC ports that are available as shown in Figure 2-61.

All installed ports
Bay 1, Card 3, Port A
Bay 1, Card 4, Port A
Bay 2, Card 3, Port A
Bay 2, Card 4, Port A
Bay 3, Card 3, Port A
Bay 3, Card 4, Port A
Bay 4, Card 3, Port A
Bay 4, Card 4, Port A

Figure 2-61 All installed ESS FC ports

We scroll to and select the adapter tree in bay 1 to limit the host FC port PC1_2 to this ESS FC port. We could also select any other port in addition to the one we have chosen. To update the Host Systems List table, we press the **Modify** button.



The restriction is now reflected in the FC Ports column, shown in Figure 2-62.

Host Systems List					
Midname	Host Type	Attachment	WWPN	Hostname/IP Address	FC Ports
PC1_1	PC Server (Win NT 4.0 or higher)	Fibre Channel attached	200000E08B00CC5B		All installed ports
PC1_2	PC Server (Win NT 4.0 or higher)	Fibre Channel attached	210000E08B017862		Bay 1, Card 3, Port A

Figure 2-62 Host Systems List with restricted host FC port PC1_2

Now we perform the configuration update, which makes the changes we made active.

Viewing changes on the Graphical View panel

Back on the Storage Allocation — Graphical View panel, we will see that the changes made are also graphically reflected. This is shown in Figure 2-63.

We see that a line is drawn from the host FC port PC1_2 to the ESS FC port three in bay one.

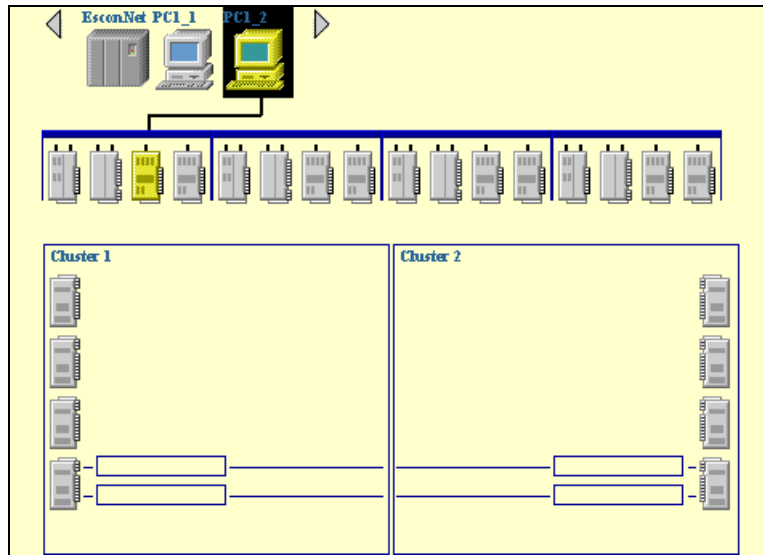


Figure 2-63 Selected host FC port, restricted to one ESS FC port

It shows us also that there are no ESS logical volumes currently assigned to the selected host FC port. This is indicated by the color of the disk group which indicates that no volume exists. This is going to change, as we are now prepared to add volumes.

2.5 Creating and assigning ESS logical volumes

As mentioned before, adding volumes to the disk groups and assigning them to the newly defined host FC ports is done in one process. It involves two steps using the Add Volume panels.

2.5.1 Adding ESS logical volumes, step one

From the Open System Storage panel, we use the **Add Volumes** button to start the definition of volumes on the newly configured disk groups, as shown in Figure 2-64.




Figure 2-64 Entry to the Add Volumes (1 of 2) panel

This takes us to the Add Volumes (1 of 2) panel, as shown in Figure 2-65.

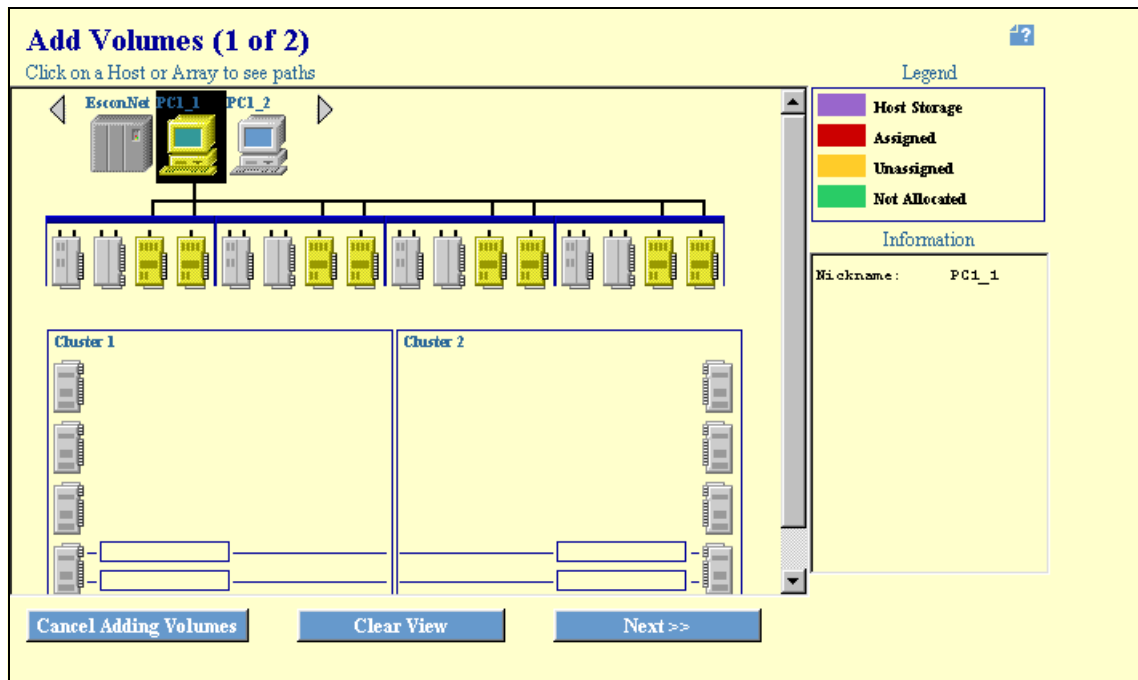


Figure 2-65 Add Volumes (1 of 2) panel with host FC port PC1_1 selected

Selecting host FC port for the volume assignment

From here, we select the host FC port **PC1_1** that we just have defined. We also have to select one of the ESS FC ports (Figure 2-66), before we can go to the next panel.




Figure 2-66 host FC port and the ESS FC port selected

Logically, there is no reason to have to select a specific ESS FC port, because we can reach the ESS logical volumes through every properly configured ESS FC port. Even in the case of the host FC port PC1_2 that we restricted to a specific ESS FC port, this would not make any difference.

Selecting the disk groups for the volume placement

By clicking on a port, we can now see the storage that we can use to define volumes on. This is shown in Figure 2-67.

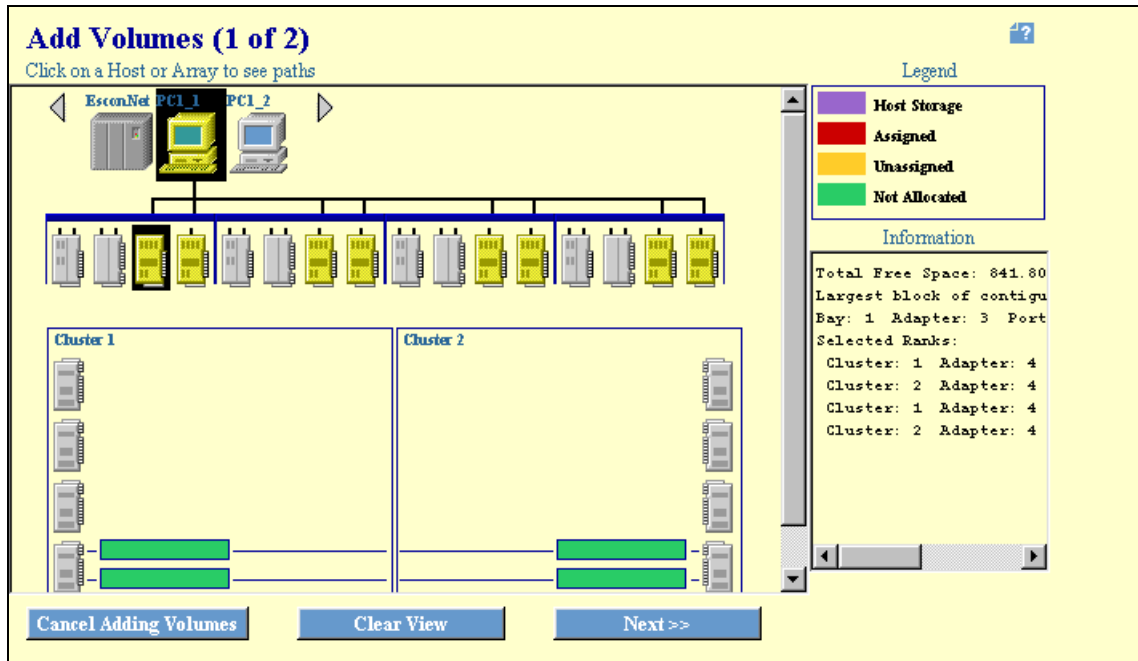


Figure 2-67 Add Volumes (1 of 2) panel with disk group status shown

There are different colors to indicate which state the disk groups are in. In our case, none of the four disk groups that we configured contains an ESS logical volume. This is indicated by a green color. Only disk groups that contain unallocated storage can be used to add volumes.

In Figure 2-67, no disk group has been manually selected by us. This means that all of them are automatically selected for the next step of the process of defining and assigning volumes. Which disk groups are selected is shown in the Information window of the panel. However, we mark the two disk groups for use in the next step by clicking on these. This highlights them with a border, as shown in Figure 2-68.




Figure 2-68 One disk group selected

2.5.2 Adding ESS logical volumes — step two

Pressing the **Next** button, we are taken to the second step of defining the volumes.

Next >>

This is the Add Volumes (2 of 2) panel, shown in Figure 2-69. This is where we define how many volumes of what size we want to place sequentially on, or spread across all prior selected disk groups.

Add Volumes (2 of 2)

Available Free Space

Storage Type	Available Capacity	Maximum Volume Size
RAID-5 Array	420.90 GB	210.45 GB
Non-RAID	0.00 GB	0.00 GB

Volume Attributes

Select a Volume Size
0 GB

Number of Volumes
[]

Storage Type
[]

Add >>

<< Remove

New Volumes

Number	Volume Size	Storage Type
Total:	0 GB	

Volume Placement

☒ Place volumes sequentially, starting in first selected storage area

☐ Spread volumes across all selected storage areas

<< Back Perform Configuration Update Cancel Configuration Update

Figure 2-69 Add Volumes (2 of 2) panel, two RAID 5 disk groups to be used

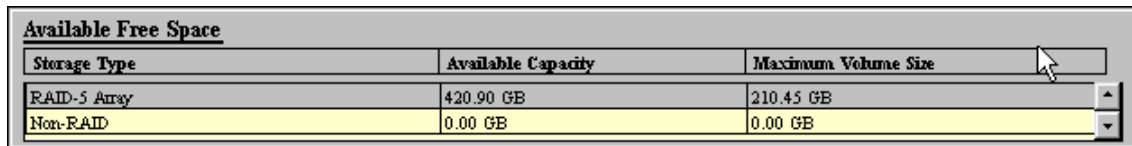
At the top of the panel, we see the Available Free Space area. We can see the available capacity and the maximum volume size of the disk groups we selected. Those numbers are in the row corresponding to the storage type we configured in the Configure Fixed Block Storage panel.

Because we did not configure and select any disk group for JBOD, the Available Capacity and Maximum Volume Size column show both an available free space of zero.

Note: An ESS logical volume can be as big as we want it to be within the boundaries of one disk group. In our case, using 36.4 GB disks, we can have a maximum volume size of 210.45 GB. Looking at the Information window on the Storage Allocation — Graphical View panel we see 210.48 GB.

Selecting the storage type

To create RAID 5 volumes, we have to click on the RAID-5 Array row shown in Figure 2-70.

The image shows a window titled "Available Free Space" containing a table with three columns: "Storage Type", "Available Capacity", and "Maximum Volume Size". There are two rows: "RAID-5 Array" and "Non-RAID". The "RAID-5 Array" row is highlighted in yellow and has a mouse cursor pointing at it. The "Non-RAID" row is also highlighted in yellow but is not selected. The table has scrollbars on the right side.

Storage Type	Available Capacity	Maximum Volume Size
RAID-5 Array	420.90 GB	210.45 GB
Non-RAID	0.00 GB	0.00 GB

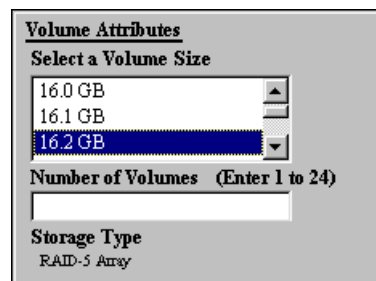
Figure 2-70 Available Free Space frame, RAID-5 storage entry selected

If no storage type has been selected, the Volume Attributes field on the left hand side shows only a volume size of zero.

Choosing the size and number of volumes

Once we have selected RAID-5 Array as the storage type, we can choose the size and the number of the ESS logical volumes that we want to configure.

We are able to choose the size of the volumes granularity, beginning from 0.1 GB up to the total capacity of our disk group, which is 210.45 GB, with a granularity of 0.1 GB. We want to define four volumes with 16.1 GB each. The reason for the volumes being exactly 16.1 GB, is that we want to all volumes the same size without wasting any space. Defining volumes with a little less or a little more space will mean that we will either fill the remaining space with smaller volumes or waste space. For example, with a volume size of 16.2 GB, we could only define 24 of them in two disk groups, which is shown in Figure 2-71.

The image shows a window titled "Volume Attributes". It contains a "Select a Volume Size" dropdown menu with three options: "16.0 GB", "16.1 GB", and "16.2 GB". The "16.2 GB" option is selected and highlighted in blue. Below the dropdown is a text field labeled "Number of Volumes (Enter 1 to 24)". At the bottom, there is a "Storage Type" label with "RAID-5 Array" displayed below it.

Volume Attributes

Select a Volume Size

16.0 GB
16.1 GB
16.2 GB

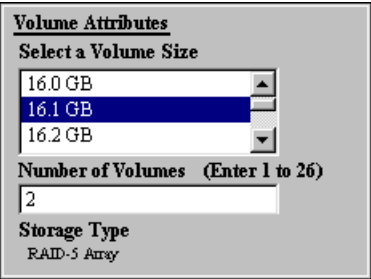
Number of Volumes (Enter 1 to 24)

Storage Type
RAID-5 Array

Figure 2-71 Volume Attributes frame, ESS logical volume size of 16.2 GB selected

Instead, we use a volume size of 16.1 GB, which allows us to place 26 of them in our selected disk groups.

To illustrate the process we only want to define two volumes here. We enter the number of volumes in the entry field, as shown in Figure 2-72.



The 'Volume Attributes' dialog box contains the following elements:

- Select a Volume Size:** A list box with three options: '16.0 GB', '16.1 GB' (highlighted in blue), and '16.2 GB'.
- Number of Volumes (Enter 1 to 26):** A text input field containing the number '2'.
- Storage Type:** A label with the text 'RAID-5 Array' below it.

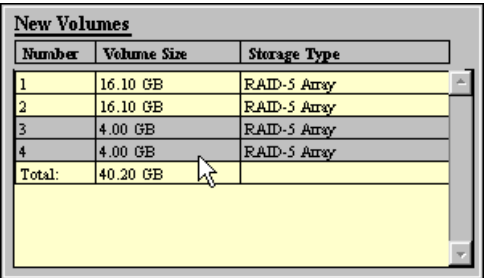
Figure 2-72 Volume Attributes frame, ESS logical volume size of 16.1 GB selected

Adding volume definitions to the New Volumes frame

We press the **Add** button to place the two volumes in the New Volumes frame.



The volumes are now shown in the New Volumes Panel, as shown in Figure 2-73. We could also add other volumes with different sizes to the list.



The 'New Volumes' frame displays a table with the following data:

Number	Volume Size	Storage Type
1	16.10 GB	RAID-5 Array
2	16.10 GB	RAID-5 Array
3	4.00 GB	RAID-5 Array
4	4.00 GB	RAID-5 Array
Total:	40.20 GB	

Below the table is a large yellow rectangular area. A mouse cursor is pointing at the 'Total' row.

Figure 2-73 New Volumes frame, two volumes selected to be removed from frame

Removing volume definitions from the New Volumes frame

To remove unwanted changes from the list, we simply highlight those volumes that we wish to remove from the list and press the **Remove** button.



Placing the volumes in selected disk groups

Now, as we have defined how many volumes and of what size we want to add, we can also decide how the selected space should be used for the placement of the volumes.

The two choices we have are shown in the Volume Placement frame. This is shown in Figure 2-74.



Figure 2-74 Volume Placement frame, spreading volumes selected

Here we can choose whether we want to place the volumes sequentially on all selected disk groups, or if we want to spread the volumes over all selected disk groups.

Placing volumes sequentially means that, in our case, the two 16.1 GB volumes would both be placed on the first selected disk group. If we wanted to configure a large number of 16.1 GB volumes, this would mean that the first 13 of them would be placed on the first disk group and the following would start to fill the second disk group.

In contrast to this, spreading volumes across all selected storage areas means exactly what it says. In our case, one of our two 16.1 GB volumes is going to be placed on each selected disk group.

Performing the configuration update

Our panel now looks like that shown in Figure 2-75, just before we apply the configuration changes by pressing the **Perform Configuration Update** button.

Add Volumes (2 of 2)

Storage Type	Available Capacity	Maximum Volume Size
RAID-5 Array	420.90 GB	210.45 GB
Non-RAID	0.00 GB	0.00 GB

Volume Attributes

Select a Volume Size

16.0 GB

16.1 GB

16.2 GB

Number of Volumes

Storage Type

RAID-5 Array

Add >>

<< Remove

New Volumes

Number	Volume Size	Storage Type
1	16.10 GB	RAID-5 Array
2	16.10 GB	RAID-5 Array
Total:	32.20 GB	

Volume Placement

☐ Place volumes sequentially, starting in first selected storage area

☒ Spread volumes across all selected storage areas

<< Back

Perform Configuration Update

Cancel Configuration Update

Figure 2-75 Add Volumes (2 of 2) panel, ready to perform the configuration update

Again, we get a window warning us about a time intensive action, and we are asked if we want to continue with the update. This is shown in Figure 2-76.

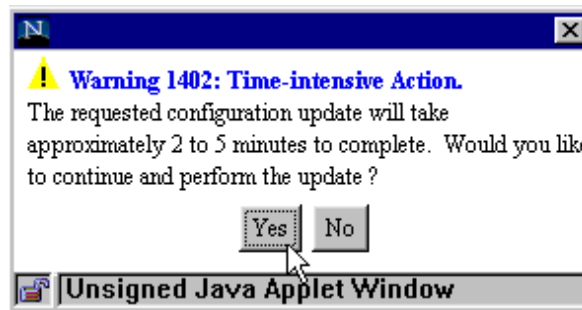


Figure 2-76 Warning window, time consuming action

As with every run of a configuration update, there is a progress window, as shown in Figure 2-77.




Figure 2-77 Progress window, volume creation

After the update has finished successfully, we are presented with the message shown in Figure 2-78.



Figure 2-78 Information window, successful volume creation

Pressing OK leads us back to the Adding Volumes (1 of 2) panel, from where we can start the process again and define other volumes.

Viewing changes on the Graphical View panel

To see the updates we have made, we press the **Storage Allocation** button. This takes us back to the Storage Allocation — Graphical View panel. There, we select the host FC port that we configured. This is shown in Figure 2-79.

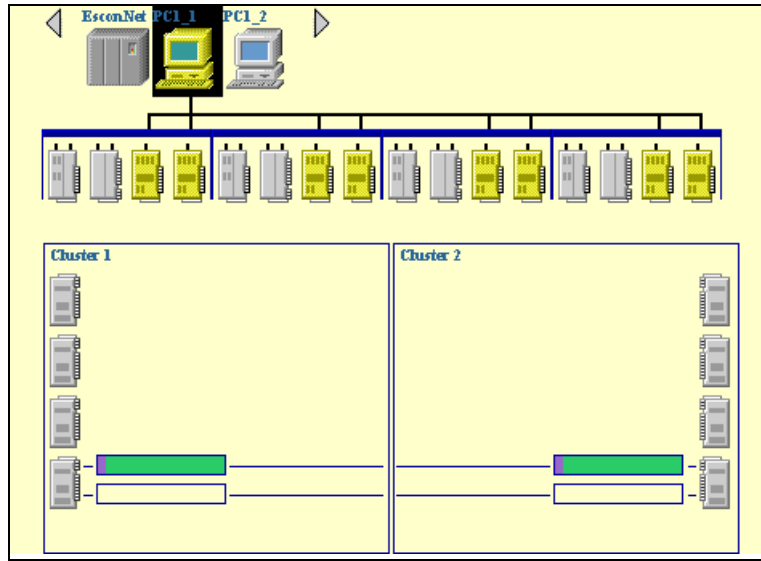


Figure 2-79 Host FC port selected

This shows the volumes defined for our host FC port. The small purple parts of both disk arrays represent the space that is filled with volumes that are assigned to the selected host FC port. The green part indicates space that is not yet allocated.

Clicking on an ESS Fibre Channel port will draw lines to the volumes configured for the selected host Fibre Channel port. This is shown in Figure 2-80.

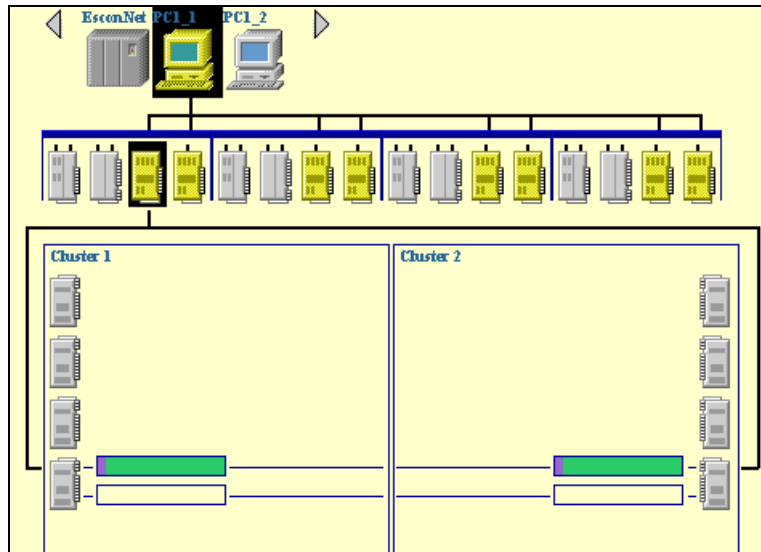


Figure 2-80 Host FC port and ESS FC port selected

The Information field now displays information relating to the configuration of the ESS FC port, which is currently undefined. This information is useful later to see which ESS FC port is configured for the appropriate topology.

We show how to configure this in 2.7, “Configuring ESS FC ports” on page 158.

After selecting one of the disk groups, the display changes again. Now the disk group has a thick border, and the Information frame has changed too. This is shown in Figure 2-81.

Information	
Cluster:	1
Adapter:	4
Loop:	A
Disk Group:	2
Storage Type:	RAID Array
Track Format:	Fixed Block (FB)
Host Storage	16.10/210.48GB (8%)
Assigned	0.00/210.48GB (0%)
Unassigned	0.00/210.48GB (0%)
Not Allocated	194.35/210.48GB (92%)

Figure 2-81 Information frame, host FC port, and disk group selected

The items that have changed on the disk group information since our last visit, are the Host Storage row and the Not Allocated row.

Host Storage is the space that now has an ESS logical volume of 16.1GB assigned to the selected host FC port.

Not Allocated means the space that can be used to define other ESS logical volumes on.

Last but not least, we can also click on the disk array without selecting a host FC port first. The panel's look changes again when only the disk group is selected. This is shown in Figure 2-82.

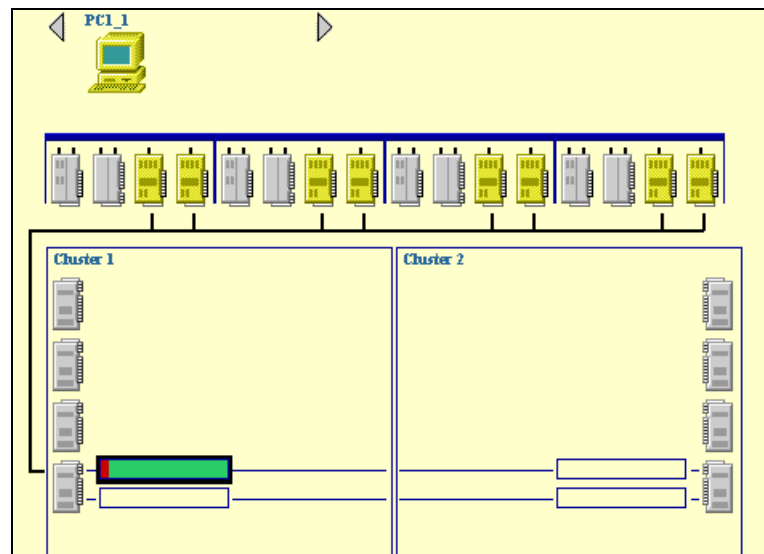


Figure 2-82 Disk group selected

To deselect everything that is currently selected, we press the **Clear View** button on the top of the panel, or we can click on the host FC port that is selected. Then we select the disk group in cluster one.

Notice that the part of the disk group that was purple before has changed to red. This is because we did not specify that we are interested in the volumes that are assigned to a particular host FC port. Instead, we now see the storage that is allocated to any of the host FC ports. The host FC port icons that have ESS logical volumes assigned to them on this disk group, now show up in the host interfaces row and are also highlighted yellow.

In our case, only the host FC port PC1_1 has a volume on this disk group assigned to it, and therefore there is only the icon of this port visible.

Again, the Information frame changed the look to represent the general information of the disk group and is still not related to a specific host FC port. The used space now shows up in the Assigned row instead of the Host Storage row. This is shown in Figure 2-83.

Information	
Cluster:	1
Adapter:	4
Loop:	A
Disk Group:	2
Storage Type:	RAID Array
Track Format:	Fixed Block (FB)
Host Storage	0.00/210.48GB (0%)
Assigned	16.10/210.48GB (8%)
Unassigned	0.00/210.48GB (0%)
Not Allocated	194.35/210.48GB (92%)

Figure 2-83 Information frame, disk group selected

Now we want to go on and define more volumes and assign them to the host FC port PC1_2.

Viewing changes on the Open System Storage panel

To perform this task, we navigate all the way back to the Add Volumes panel. From the Storage Allocation — Graphical View panel, we press the **Open System Storage** button.



This takes us to the Open System Storage panel, as shown in Figure 2-84.

Open System Storage

Host Systems

Nickname	Host Type	Attachment	WWPN	Hostname/IP Address
PC1_1	PC Server (Win NT 4.0 or higher)	FC	200000E08B00CC5B	
PC1_2	PC Server (Win NT 4.0 or higher)	FC	210000E08B017862	

Assigned Volumes (Total: 2 volumes)

Volume	Vol Type	Size	Storage Type	Location	LSS	Shared
700-18486	Open System	16.1 GB	RAID Array	Device Adapter Pair 4 Cluster 2, Loop A Array 1, Vol 000	LSS: 017	No
600-18486	Open System	16.1 GB	RAID Array	Device Adapter Pair 4 Cluster 1, Loop A Array 2, Vol 000	LSS: 016	No

Modify Host Systems
Configure Host Adapter Ports
Configure Disk Groups

Add Volumes
Modify Volume Assignments

Figure 2-84 FC port PC1_1 with two assigned volumes

Now when we select the host FC port **PC1_1** in the Host Systems field, the associated volumes appear in the Assigned Volumes field with all the relevant details.

For host FC port PC1_2, the Assigned Volumes table is still empty.

Creating two more ESS logical volumes

Now we do the same as we did before with the host FC port PC1_1, for PC1_2. We are going to create and assign two volumes of 64 GB, one of them in disk group 1 and disk group 2 in loop B of the device adapter pair 4. This is done exactly like we did before with the host FC port PC1_1, so we only show the results.

Viewing changes on the Graphical View panel

Back on the Storage Allocation — Graphical View panel we see how the panel now looks using the View All Storage mode, shown in Figure 2-85.

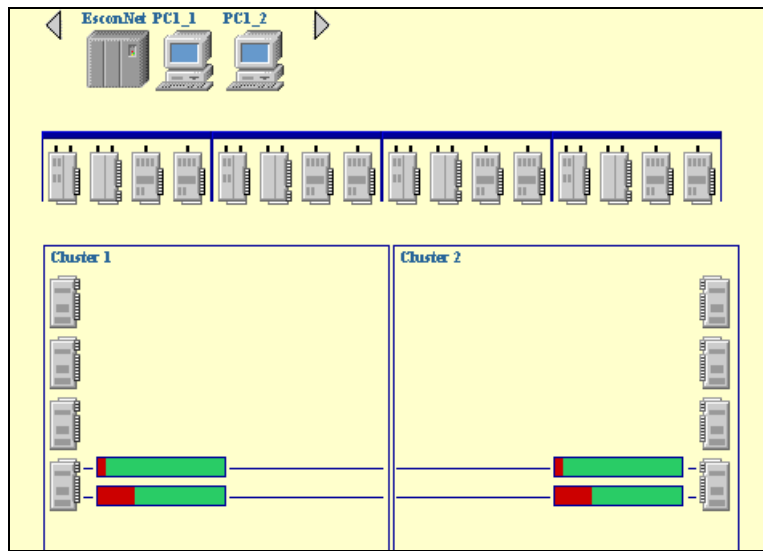


Figure 2-85 View All Storage mode

The four areas in the four disk groups now represent the space that is used for ESS logical volumes. We cannot see to which host FC port they are assigned. To discover what volumes are assigned to PC1_2, we select this host FC port. At this stage, we see that the two disk groups have purple spaces on them, which means they are assigned to the selected host FC port.

We select the host PC1_2, and we see that lines now appear connecting the host port, ESS FC port, and storage volumes. This is shown in Figure 2-86.

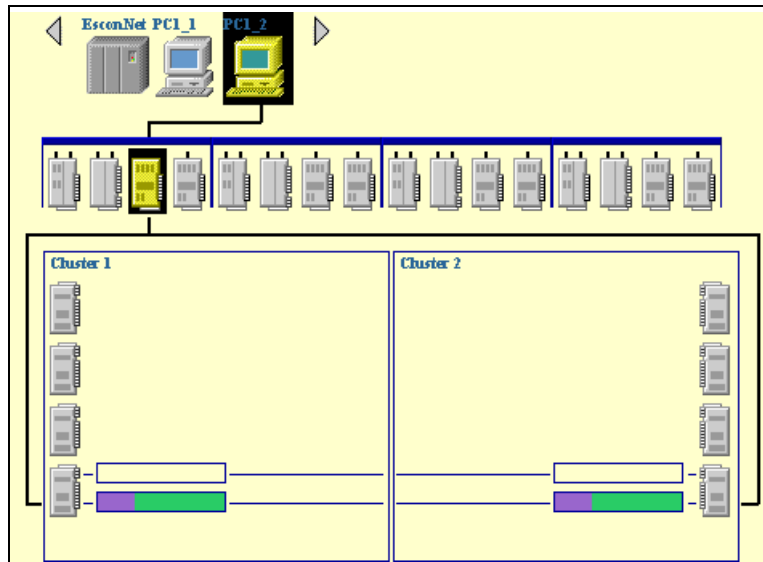


Figure 2-86 Host FC port and ESS FC port selected

Viewing changes on the Tabular View panel

Now we will look at the Storage Allocation — Tabular View panel. We reach this from here using the **Tabular View** button on the top right of the panel.

Tabular View

This shows detailed information on all host FC ports assigned ESS logical volumes in one table, as shown in Figure 2-87.

Storage Allocation -- Tabular View							
List of Assigned Volumes				Print Table	Perform Sort	Graphical View	
no sort	no sort	no sort	no sort	no sort	no sort	no sort	no sort
Host/SSID	LSS/LCU	Volume	Type	Size	Host Adapter	Location	Shared
PC1_1	LSS: 016	600-18486	Open System	016.1 GB	Fibre Channel FC Port A ID 00, LUN 0000	Device Adapter Pair 4 Cluster 1, Loop A Array 2, Vol 000	No
PC1_1	LSS: 017	700-18486	Open System	016.1 GB	Fibre Channel FC Port A ID 00, LUN 0001	Device Adapter Pair 4 Cluster 2, Loop A Array 1, Vol 000	No
PC1_2	LSS: 016	601-18486	Open System	064.0 GB	Fibre Channel FC Port A ID 00, LUN 0000	Device Adapter Pair 4 Cluster 1, Loop B Array 2, Vol 001	No
PC1_2	LSS: 017	701-18486	Open System	064.0 GB	Fibre Channel FC Port A ID 00, LUN 0001	Device Adapter Pair 4 Cluster 2, Loop B Array 1, Vol 001	No

Figure 2-87 Displays of all four volume assignments

We see all the information about the location and the current assignments here, as well as the ESS logical volume ID. This is the ID that will help us later to identify the ESS logical volumes on our host system. Right now, there are not too many entries in the table.

Later on, as we define more volumes, the table will become large. To make it easier to view the table, we can also perform a sort on the table. The sort can be performed using three choices of levels using the drop down list boxes on the top of each column. To perform the sort, we press the **Perform Sort** button.

Now, we will modify ESS logical volume assignments.

2.6 Modifying ESS logical volume assignments

At this time we would only have to configure all the ESS FC ports to enable the ESS to present volumes to a Fibre Channel host. There are further changes of the assignment of volumes to ports that might have to be performed. For this, we use the Modify Volume Assignments panel.

For example, we can share volumes between host FC ports. These ports can be in the same host which uses, for example, the Subsystem Device Driver (SDD), to increase the bandwidth and availability, or in different hosts, for example, for clustering. Or, we might want to unassign volumes that were shared before, or we might want to unassign volumes to keep them unused for later reuse. We cover all this in the following topics.

2.6.1 Sharing ESS logical volumes between two host FC ports

We click on the **Modify Volume Assignments** button as shown in Figure 2-88.



Figure 2-88 Entry point for modifying volume assignments

This takes us to the Modify Volume Assignments panel, as shown in Figure 2-89.

Modify Volume Assignments

[Volume Assignments](#)
[Refresh Status](#)
[Print Table](#)
[Perform Sort](#)

no sort

no sort

no sort

no sort

no sort

no sort

no sort

no sort

Volume	Location	ESS	Volume Type	Size	Storage Type	Host Port	Host Nidnames
600-18486	Device Adapter Pair 4 Cluster 1, Loop A Array 2, Vol 000	16	Open System	016.1 GB	RAID Array	Fibre Channel ID 00, LUN 0000	PC1_1
601-18486	Device Adapter Pair 4 Cluster 1, Loop B Array 2, Vol 001	16	Open System	064.0 GB	RAID Array	Fibre Channel ID 00, LUN 0000	PC1_2
700-18486	Device Adapter Pair 4 Cluster 2, Loop A Array 1, Vol 000	17	Open System	016.1 GB	RAID Array	Fibre Channel ID 00, LUN 0001	PC1_1
701-18486	Device Adapter Pair 4 Cluster 2, Loop B	17	Open System	064.0 GB	RAID Array	Fibre Channel ID 00, LUN 0001	PC1_2

[Action](#)

☐ Assign selected volume(s) to target hosts
 ☐ Use same ID/Lun in source and target
 ☐ Unassign selected volume(s) from target hosts

[Target Hosts](#)

[Perform Configuration Update](#)
[Cancel Configuration Update](#)

Figure 2-89 Four volume assignments, one for each existing volume

The Volume Assignments table provides a comprehensive overview of all configured volume associations. We can also sort the table, specifying the column which we want to use as sort criteria. The rows do not represent the volumes. Instead there will be a row for each association that a volume has as we will see later on.

There is an Action field, which is grayed out, and a Target Hosts field, which is empty, when no volume association is selected. Now we show their interaction.

Figure 2-90 shows the two assignments of the 64 GB volumes assigned to host FC port PC1_2. Both assignments are selected and ready for modification.

The screenshot shows the 'Volume Assignments' interface. At the top, there are buttons for 'Refresh Status', 'Print Table', and 'Perform Sort'. Below these are several dropdown menus for sorting by different columns, all set to 'no sort'. The main table has columns: Volume, Location, LSS, Volume Type, Size, Storage Type, Host Port, and Host Nicknames. Two rows are selected, highlighted in yellow. Below the table is an 'Action' field with three radio buttons: 'Assign selected volume(s) to target hosts' (selected), 'Use same ID/Lun in source and target' (disabled), and 'Unassign selected volume(s) from target hosts' (disabled). To the right of the Action field is a 'Target Hosts' field, which is empty. At the bottom are buttons for 'Perform Configuration Update' and 'Cancel Configuration Update'.

Volume	Location	LSS	Volume Type	Size	Storage Type	Host Port	Host Nicknames
	Cluster 2, Loop A Array 1, Vol 000					ID 00, LUN 0001	
600-18486	Device Adapter Pair 4 Cluster 1, Loop A Array 2, Vol 000	16	Open System	016.1 GB	RAID Array	Fibre Channel ID 00, LUN 0000	PC1_1
601-18486	Device Adapter Pair 4 Cluster 1, Loop B Array 2, Vol 001	16	Open System	064.0 GB	RAID Array	Fibre Channel ID 00, LUN 0000	PC1_2
701-18486	Device Adapter Pair 4 Cluster 2, Loop B Array 1, Vol 001	17	Open System	064.0 GB	RAID Array	Fibre Channel ID 00, LUN 0001	PC1_2

Action

☒ Assign selected volume(s) to target hosts
☐ Use same ID/Lun in source and target
☐ Unassign selected volume(s) from target hosts

Target Hosts

Perform Configuration Update Cancel Configuration Update

Figure 2-90 Two selected volume assignments, Action field accessible

Selecting ESS logical volume assignments for modification

Selections of more than one assignment can be done in two ways.

The first case is if the assignments to be selected are not adjacent. Then we click on the first assignment to be selected, press and hold the Ctrl key, and click on the next assignment to be selected until we have all the assignments we want selected. While holding down the Ctrl key, we can also click on an already selected assignment to deselect it.

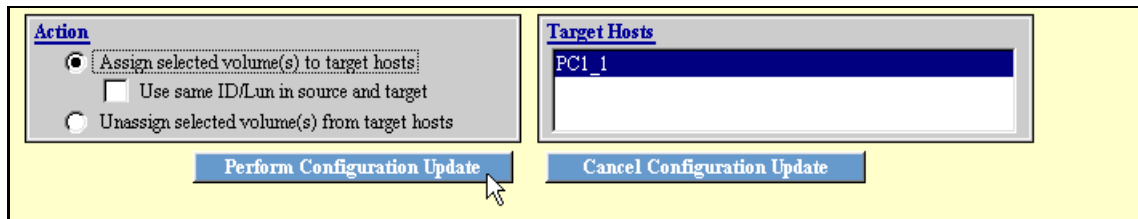
The second case is if the assignments to be selected are adjacent. Of course to avoid the application of the first selection method, we can sort the assignments to show up adjacent. We sort the assignment here by its host FC port. For adjacent assignments, we click on the first assignment to be selected, press and hold the Shift key and click on the last assignment to be selected. This selects all assignments in between.

If, for some reason, one of the assignments just selected using the Shift key should not be modified, we are also able to deselect this assignment by pressing the Ctrl key and clicking on it. In the same way, we can also add other particular assignments to the selection.

Once we have done this, we go on and choose what we want to do with the assignments using the Action field.

Creating new volume assignments for selected host FC port

With the radio button, we decide if we want to create a new volume assignment using the volume associated with selected assignments, or remove both selected volume assignments. We want to create new assignments to share the volumes and so we select the corresponding radio button shown in Figure 2-91.



The screenshot shows a configuration window with two main sections. The 'Action' section on the left contains two radio buttons: 'Assign selected volume(s) to target hosts' (which is selected) and 'Unassign selected volume(s) from target hosts'. Below these is a checkbox labeled 'Use same ID/Lun in source and target'. The 'Target Hosts' section on the right is a list box containing the text 'PC1_1'. At the bottom of the window are two buttons: 'Perform Configuration Update' and 'Cancel Configuration Update'. A mouse cursor is pointing at the 'Perform Configuration Update' button.

Figure 2-91 Adding assignments to other available host FC ports selected

Selecting this allows us to choose the host FC port from those displayed in the Target Hosts field. We have only the host FC port **PC1_1** available for this action. If there were to be more host FC ports configured, we would have to choose and select the one on which we want the assignments altered, by highlighting its name.

Performing the configuration update

To apply the changes, we press the **Perform Configuration Update** button. The progress bar indicates that the configuration update is taking place, as shown in Figure 2-92.

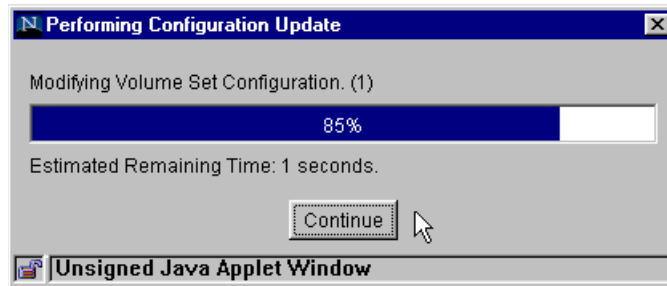


Figure 2-92 Progress window, creating volume assignments

If successful, we get the message shown in Figure 2-93.




Figure 2-93 Volume assignments successfully created for specified host FC ports

If we press OK, we do not change back to the Open System Storage panel. Instead, we see the new volume assignments to the host FC port PC1_1, as shown in Figure 2-94.

Modify Volume Assignments ?

Volume Assignments
Refresh Status
Print Table
Perform Sort

no sort ▼
no sort ▼
no sort ▼
no sort ▼
second ▼
no sort ▼
no sort ▼
first ▼

Volume	Location	LSS	Volume Type	Size	Storage Type	Host Port	Host Nicknames
600-18486	Device Adapter Pair 4 Cluster 1, Loop A Array 2, Vol 000	16	Open System	016.1 GB	RAID Array	Fibre Channel ID 00, LUN 0000	PC1_1
700-18486	Device Adapter Pair 4 Cluster 2, Loop A Array 1, Vol 000	17	Open System	016.1 GB	RAID Array	Fibre Channel ID 00, LUN 0001	PC1_1
701-18486	Device Adapter Pair 4 Cluster 2, Loop B Array 1, Vol 001	17	Open System	064.0 GB	RAID Array	Fibre Channel ID 00, LUN 0003	PC1_1
601-18486	Device Adapter Pair 4 Cluster 1, Loop B	16	Open System	064.0 GB	RAID Array	Fibre Channel ID 00, LUN 0002	PC1_1

Action

☐ Assign selected volume(s) to target hosts

☐ Use same ID/Lun in source and target

☐ Unassign selected volume(s) from target hosts

Target Hosts

Perform Configuration Update
Cancel Configuration Update

Figure 2-94 Assignments for the host FC port PC1_1

Viewing changes on the Graphical View panel

Back on the Storage Allocation — Graphical View panel, we now see that the host FC port PC1_1 has volumes assigned to it on all four disk groups as shown in Figure 2-95.

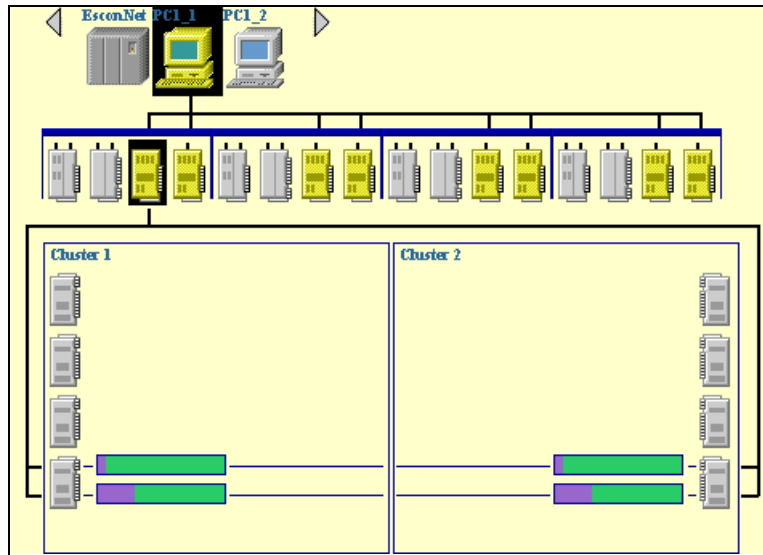


Figure 2-95 PC1_1 and ESS FC port selected

In fact there is no volume configured on the ESS which is not assigned to this port. This is indicated by the lack of any colored space on the disk group, apart from the color of purple for the selected host FC port assigned volumes, or green for unused space.

When a host FC port is selected, there is no indication that two of the volumes are shared between both the defined host FC ports. To see if ESS logical volumes are shared, we have to select a disk group. In Figure 2-96, we see a disk group selected, that contains assigned volumes.

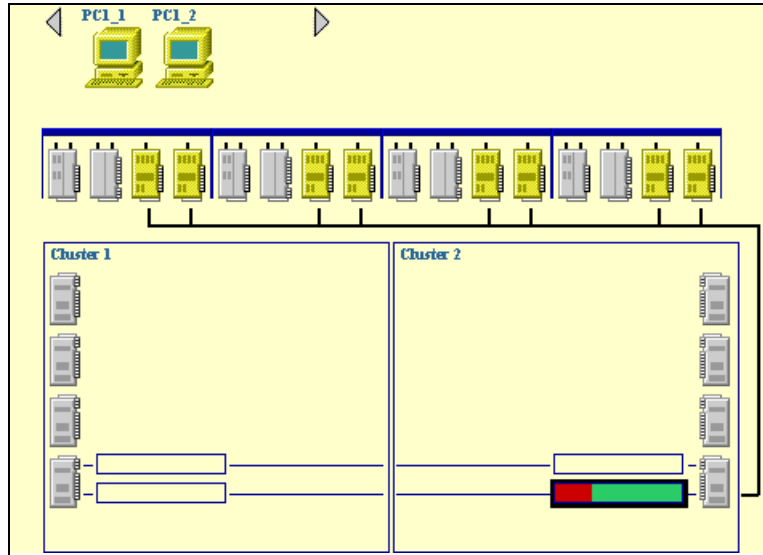


Figure 2-96 Disk group containing shared ESS logical volume(s) selected

The red color indicates that it contains volumes that are assigned to the host FC ports that show up highlighted yellow in the host interfaces row. When we compare this with Figure 2-95, we see that the amount of space used for this disk group is the same. Also, when we select the host FC port **PCI_2**, we will see the same amount of storage assigned to it, which tells us that the volumes must be shared between the two.

Viewing changes on the Tabular View panel

If we click on the **Tabular View** button and get the Storage Allocation — Tabular View panel as shown in Figure 2-97, we see two rows for each of the two 64 GB sized volumes. Once for the host FC port **PCI_1** and once for **PCI_2**. Again, the table does not represent the volumes but the assignments the volumes have, if any. The rightmost column indicates the sharing of the volumes as well.

Storage Allocation -- Tabular View							
List of Assigned Volumes				Print Table	Perform Sort	Graphical View	
no sort	no sort	no sort	no sort	no sort	no sort	no sort	no sort
Host/SSID	LSS/LCU	Volume	Type	Size	Host Adapter	Location	Shared
PC1_1	LSS: 016	600-18486	Open System	016.1 GB	Fibre Channel FC Port A ID 00, LUN 0000	Device Adapter Pair 4 Cluster 1, Loop A Array 2, Vol 000	No
PC1_1	LSS: 017	700-18486	Open System	016.1 GB	Fibre Channel FC Port A ID 00, LUN 0001	Device Adapter Pair 4 Cluster 2, Loop A Array 1, Vol 000	No
PC1_1	LSS: 016	601-18486	Open System	064.0 GB	Fibre Channel FC Port A ID 00, LUN 0002	Device Adapter Pair 4 Cluster 1, Loop B Array 2, Vol 001	Yes
PC1_1	LSS: 017	701-18486	Open System	064.0 GB	Fibre Channel FC Port A ID 00, LUN 0003	Device Adapter Pair 4 Cluster 2, Loop B Array 1, Vol 001	Yes
PC1_2	LSS: 016	601-18486	Open System	064.0 GB	Bay 1, Card 3 FC Port A ID 00, LUN 0000	Device Adapter Pair 4 Cluster 1, Loop B Array 2, Vol 001	Yes
PC1_2	LSS: 017	701-18486	Open System	064.0 GB	Bay 1, Card 3 FC Port A ID 00, LUN 0001	Device Adapter Pair 4 Cluster 2, Loop B Array 1, Vol 001	Yes

Figure 2-97 Tabular View, two volumes shared, two not shared

Converting volume assignments table in HTML format

To view the table offline, or to print it, this panel features the **Print Table** button.

Print Table

Then a window as shown in Figure 2-98 will appear, advising that the information is going to be displayed in another browser window.

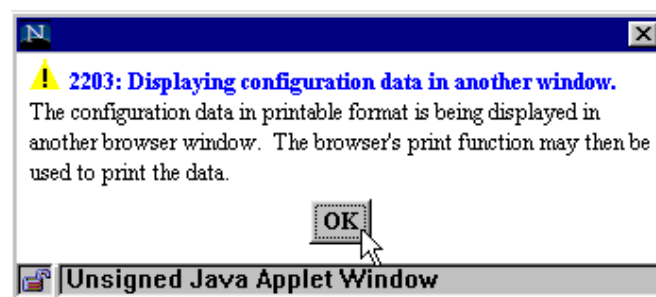


Figure 2-98 Information window, opening another browser window.

Now the content of the table is displayed in HTML format, as shown in Figure 2-99.

Host/SSID	LSS/LCU	Volume	Type	Size	Host Adapter	Location	Shared
PC1_1	LSS: 016	600-18486	Open System	016.1 GB	Fibre Channel FC Port A ID 00, LUN 0000	Device Adapter Pair 4 Cluster 1, Loop A Array 2, Vol000	No
PC1_1	LSS: 017	700-18486	Open System	016.1 GB	Fibre Channel FC Port A ID 00, LUN 0001	Device Adapter Pair 4 Cluster 2, Loop A Array 1, Vol000	No
PC1_1	LSS: 016	601-18486	Open System	064.0 GB	Fibre Channel FC Port A ID 00, LUN 0002	Device Adapter Pair 4 Cluster 1, Loop B Array 2, Vol001	Yes
PC1_1	LSS: 017	701-18486	Open System	064.0 GB	Fibre Channel FC Port A ID 00, LUN 0003	Device Adapter Pair 4 Cluster 2, Loop B Array 1, Vol001	Yes
PC1_2	LSS: 016	601-18486	Open System	064.0 GB	Bay 1, Card 3 FC Port A ID 00, LUN 0000	Device Adapter Pair 4 Cluster 1, Loop B Array 2, Vol001	Yes
PC1_2	LSS: 017	701-18486	Open System	064.0 GB	Bay 1, Card 3 FC Port A ID 00, LUN 0001	Device Adapter Pair 4 Cluster 2, Loop B Array 1, Vol001	Yes

Figure 2-99 HTML file, List of Assigned Volumes table, list of existing assignments

It can now be printed from within the browser or it can also be saved as an HTML file for later viewing.

2.6.2 Removing volume assignments from host FC ports

The second configuration change, regarding the volume assignments we want to show, is removing the assignments of the ESS logical volume in disk group two of loop B of the SSA device adapter pair 4, from both host FC ports.

Again, we click on the **Modify Volume Assignments** button, as shown in Figure 2-100.



Figure 2-100 Entry point for modifying volume assignments

Removing the assignments of one ESS logical volume

First as we did in, “Selecting ESS logical volume assignments for modification” on page 144, we select the volume assignments for both host FC ports. Then, when the Action field becomes available, we select the radio button to **Unassign selected volume(s) from target hosts**, which means, remove selected ESS logical volume assignments for the selected host FC ports. The Target Hosts field now lists the host FC ports to which the assignments belong to. Here we select both host FC ports to remove the assignments. This is shown in Figure 2-101.




Figure 2-101 Removing volume assignments from both host FC ports.

Doing so makes the warning message shown in Figure 2-102 appear, stating that this action will cause ESS logical volumes to become isolated for the selected host FC port.




Figure 2-102 Warning window one, volumes to be isolated from host FC ports

After pressing the **Perform Configuration Update** button, a second warning window will show up as shown in Figure 2-103. This time, the window gives us the choice to abort the configuration update.

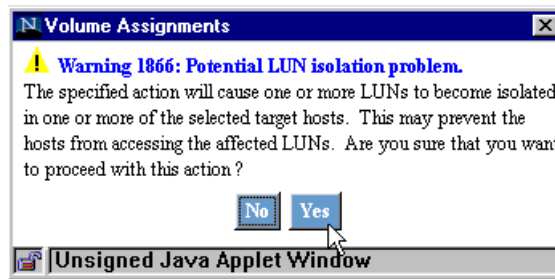


Figure 2-103 Warning window two, volumes to be isolated from host FC ports

We press the **Yes** button to proceed with the update and get a progress bar and a message window as shown in Figure 2-104 and, as we already know from the other actions we have done using this panel.

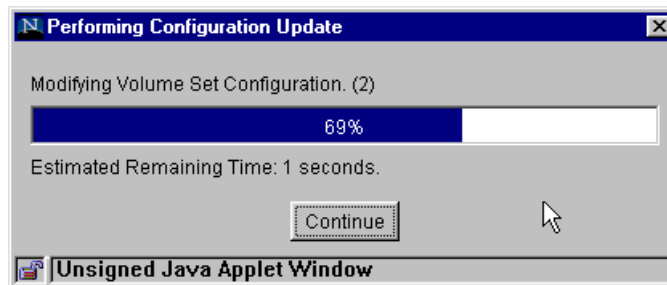


Figure 2-104 Progress window, applying assignment changes

In Figure 2-105 we are told that the volume assignments have been successfully removed.




Figure 2-105 Information window, volume assignments successfully removed

Now the Volume Assignments table shows the just unassigned volume. This table represents the assignments made for ESS logical volumes, not the volumes themselves. An exception is if volumes are not assigned to any host FC port. Then there will be one placeholder row that can be used to create an assignment. This row selected is shown in Figure 2-106.

Viewing changes on the Graphical View panel

Now we want to see the changes on the Storage Allocation — Graphical View panel. Figure 2-108 and Figure 2-109 show the selected host FC ports PC1_1 and PC1_2 with the disk groups that contain volumes that are assigned to it.

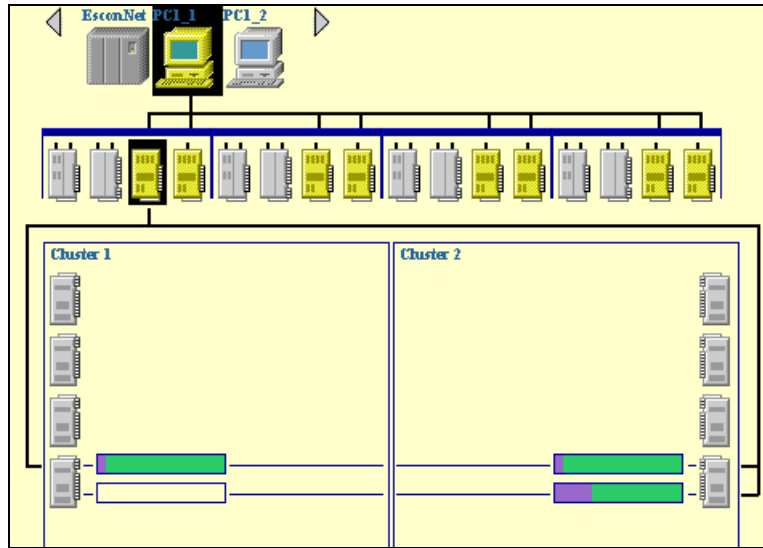


Figure 2-108 No volume assignment to PC1_1 in disk group 2 of cluster 1 loop B

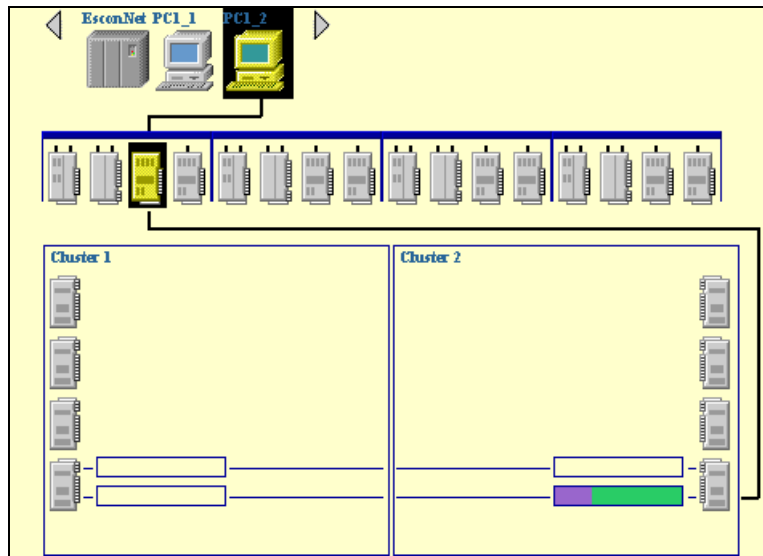


Figure 2-109 No volume assignment to PC1_2 in disk group 2 of cluster 1 loop B

The space of the disk groups that contains ESS logical volumes that are assigned to the selected host FC port show up in purple, and unused space in green.

Both host FC ports have no ESS logical volumes assigned to them in the disk group in cluster 1, adapter pair four, loop B. This is indicated by the absence of colors.

However, on the particular disk group space is allocated for volumes. This means that there is an existing volume (or more), which is not yet assigned to any host FC port. To make this allocated space visible, we press the **View All Storage** button and see all the allocated space on the disk groups. This is shown in Figure 2-110.

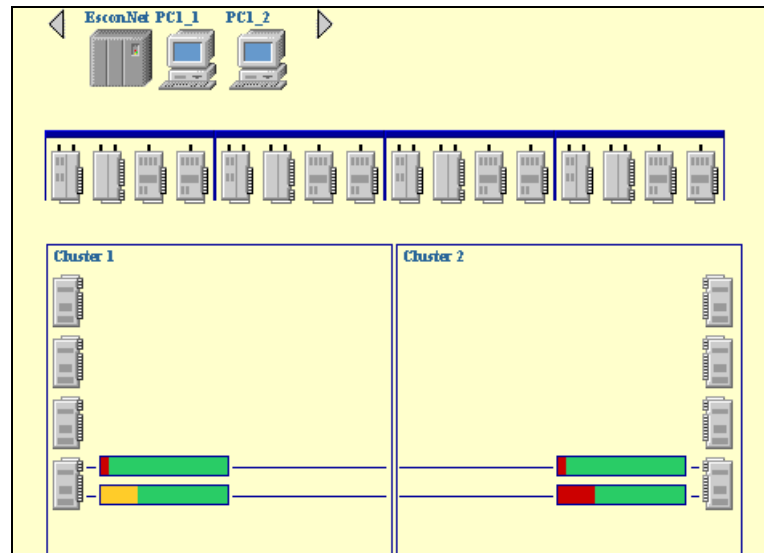


Figure 2-110 Unassigned volume(s) in disk group 2 of cluster 1 loop B

Allocated, but unassigned, volumes are indicated by the yellow colored space in the disk group. The space in the other disk groups, that contains allocated and assigned volumes shows up in the color red.

Another way to find out about the state of the space in a disk group is to select it. In our example, shown in Figure 2-111, we select the disk group in which we have an unassigned volume.

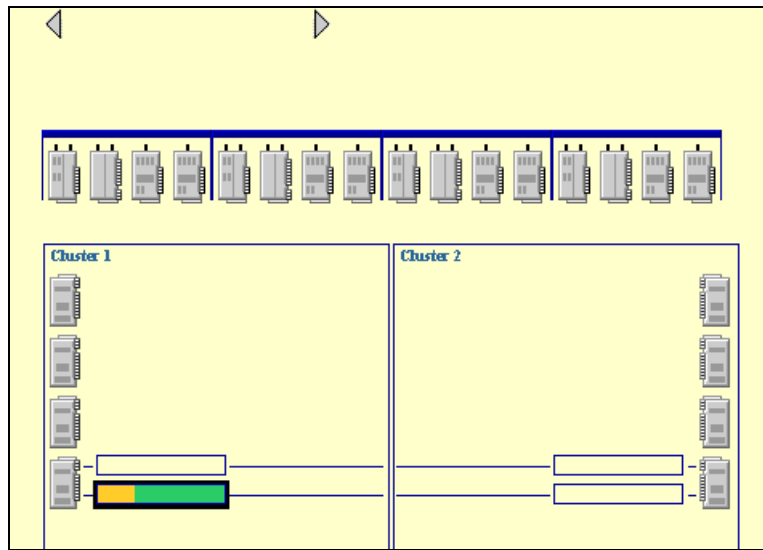


Figure 2-111 Disk group selected, unassigned volume(s) in selected disk group

Again, the space used for unassigned volumes is displayed in yellow. Also, we see in the host interfaces row the interfaces that have volumes assigned to them in the selected disk group, which is empty.

Viewing Changes on the Tabular View

Pressing the tabular **View** button from within the Storage Allocation — Graphical View panel, displays the table containing all existing volume assignments, as shown in Figure 2-112.

Storage Allocation -- Tabular View							
List of Assigned Volumes				Print Table	Perform Sort	Graphical View	
no sort	no sort	no sort	no sort	no sort	no sort	no sort	no sort
Host/SSID	LSS/LCU	Volume	Type	Size	Host Adapter	Location	Shared
PC1_1	LSS: 016	600-18486	Open System	016.1 GB	Fibre Channel FC Port A ID 00, LUN 0000	Device Adapter Pair 4 Cluster 1, Loop A Array 2, Vol 000	No
PC1_1	LSS: 017	700-18486	Open System	016.1 GB	Fibre Channel FC Port A ID 00, LUN 0001	Device Adapter Pair 4 Cluster 2, Loop A Array 1, Vol 000	No
PC1_1	LSS: 017	701-18486	Open System	064.0 GB	Fibre Channel FC Port A ID 00, LUN 0003	Device Adapter Pair 4 Cluster 2, Loop B Array 1, Vol 001	Yes
PC1_2	LSS: 017	701-18486	Open System	064.0 GB	Bay 1, Card 3 FC Port A ID 00, LUN 0001	Device Adapter Pair 4 Cluster 2, Loop B Array 1, Vol 001	Yes

Figure 2-112 Tabular View, four volume assignments, three volumes

Here we see that there is no longer any assignment for the volume we unassigned. Unassigned volumes only show up in the Volume Assignments table in the Modify Volume Assignments panel.

Now, as we have ESS logical volumes assigned to host FC ports, and we have also made adjustments to the configuration, there is only one step left to be performed.

2.7 Configuring ESS FC ports

The only configuration part left to enable us to access the volumes is the configuration of the ESS FC ports for the topology that the ports are going to connect to. This configuration step does not depend on any other configuration, so we left it until we had finished the volume configurations.

As we already know, the selection of an entity in the Storage Allocation — Graphical View panel changes the information in the Information frame of the panel. Clicking on an icon of an ESS FC port in the ESS interfaces row displays information about the adapter. There we see the physical location in the ESS, the topology the port is configured for and the WWPN. For any ESS FC port that we have not configured, the window looks like that shown in Figure 2-113.

Information	
Adapter Type:	FC host attachment
IO/Bay:	1
Adapter:	3
Port:	A
WWPN:	100000000C9251126
Topology:	Undefined

Figure 2-113 Information frame, ESS FC port topology undefined

2.7.1 Defining the topology for ESS FC ports

Pressing the **Configure Host Adapter Ports** button takes us to the Configure Host Adapter Ports panel as shown in Figure 2-114.

Modify Host Systems	Configure Host Adapter Ports	Configure Disk Groups
Add Volumes	Modify Volume Assignments	

Figure 2-114 Entry point for configuring ESS FC ports

Again, the naming convention in the ESS Specialist can be a little misleading. The ESS has built in Fibre Channel adapters which provides the connectivity to the SAN. An ESS, in its simplest form, is a Fibre Channel attached storage device, not a host. We will refer to such a device as an ESS FC port to avoid any confusion.

Selecting the interface port to be configured

The Configure Host Adapter Ports panel is shown in Figure 2-115.

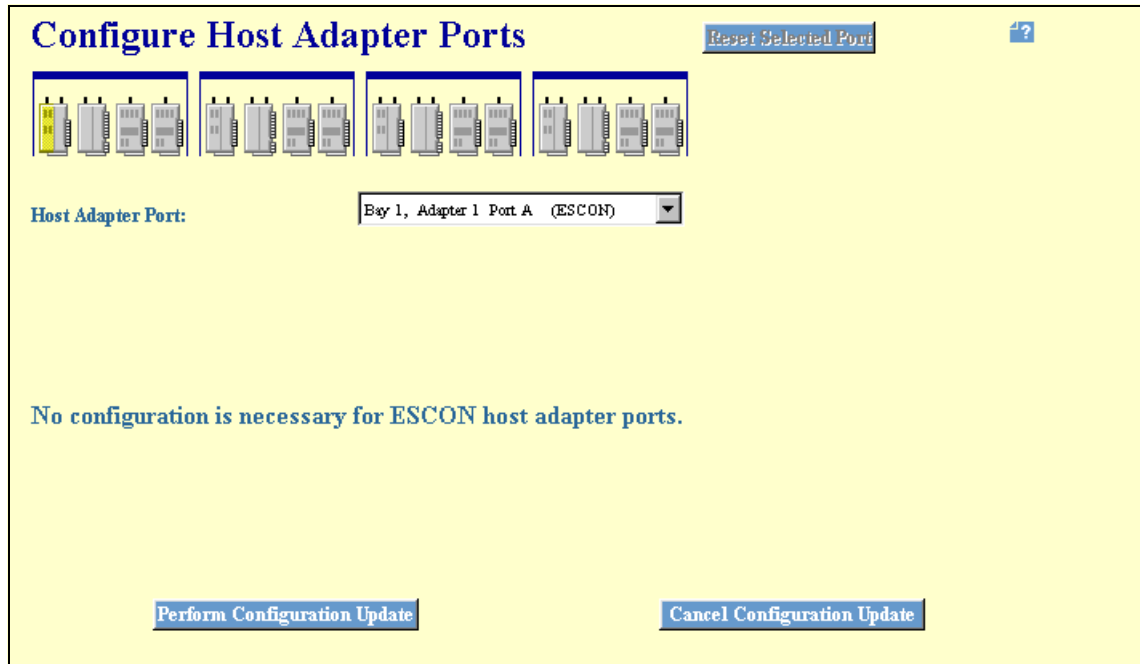


Figure 2-115 Configure Host Adapter Ports panel, ESCON port selected

From this panel, we select the port of the ESS interface adapters to be configured, depending on the type of the adapter. We can select the ports by clicking on the icons of the adapter (left or right on any adapter with two ports), or by selecting the port from the drop down list as shown in Figure 2-116.

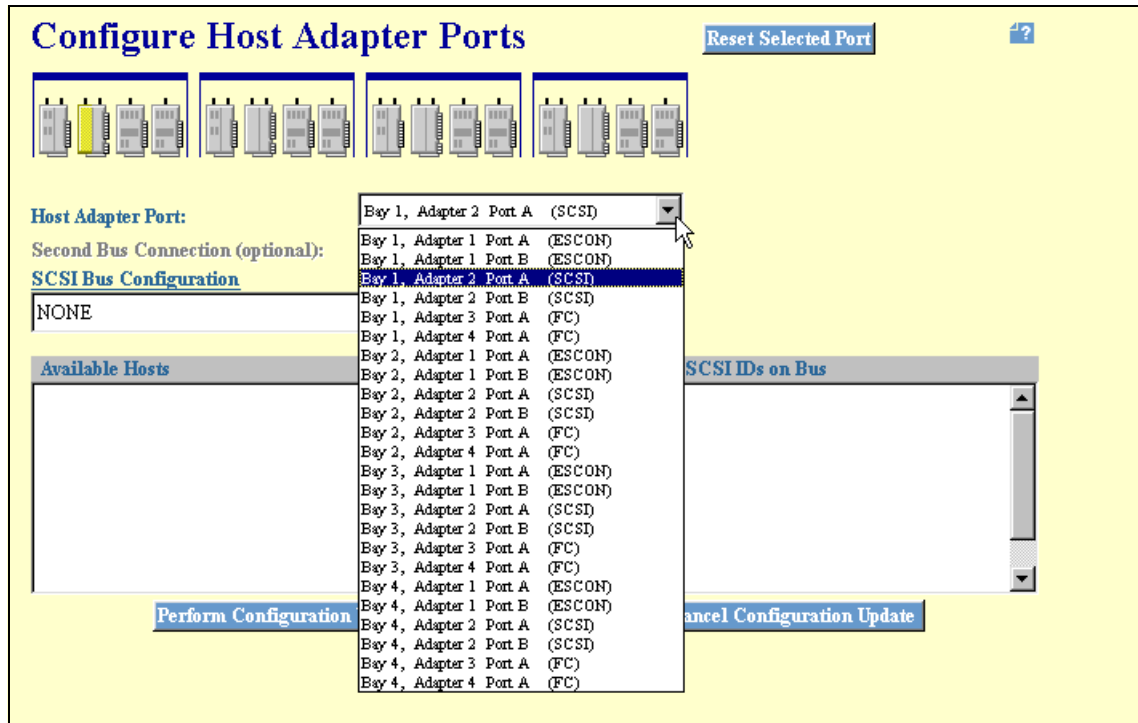


Figure 2-116 Port selection drop down list, SCSI port selected

There is no configuration necessary for ESCON. There are different configuration options for SCSI and Fibre Channel. However, we will focus on the FC part.

By selecting an FC port, we are presented with the panel, as shown in Figure 2-117.

Configure Host Adapter Ports Reset Selected Port ?

Host Adapter Port: Bay 1, Adapter 3 Port A (FC)

Storage Server Attributes

Access_Restricted Fibre Channel Access Mode

FC Port Attributes

10000000C9251126 World-Wide Port-Name

Undefined Fibre Channel Topology

Perform Configuration Update Cancel Configuration Update

Figure 2-117 Configure Host Adapter Ports panel, FC port selected

The port we used is in bay 1, adapter 3, and is port A. It is called 'A' even if it is the only port on the adapter.

Understanding the Storage Server Attributes field

As shown in Figure 2-118, the Storage Server Attributes field — also known as the logical unit number (LUN) access mode — which specifies the channel access mode, can only be changed by an IBM SSR using the service terminal.

Storage Server Attributes

Access_Restricted Fibre Channel Access Mode

Figure 2-118 Storage Server Attributes field, ESS configured for Access_Restricted

There are two Fibre Channel Access Modes to which the ESS can be configured:

► Access_Any

In Access_Any mode, any host FC port that is not defined in the ESS Specialist, can access all non-AS/400, ESS logical volumes. In Access-Any mode, the ESS Specialist displays an anonymous pseudo-host FC port icon, which is shown in Figure 2-119 and which is taken from another ESS, because the one we used here is configured for Access_Restricted. An anonymous host FC port is not a real host FC port in disguise and connected to the storage server. It represents all host FC ports (if any) that are connected to the ESS and that are *not* configured on the ESS. The anonymous host FC ports do not have an access profile defined. You can access specific ESS logical volumes that are defined in the ESS, by host FC ports that are not identified by the ESS.

► Access_Restricted

In Access_Restricted mode, a connected host FC port, which is not configured with the WWPN, cannot access any ESS logical volume. When this host FC port gets configured by using the WWPN, it will only see the volumes for which it is configured. Once a host FC port is configured, there are no differences.




Figure 2-119 Anonymous host FC port icon in an ESS configured for Access_Any

Understanding access profiles

Whichever access mode is chosen, any host FC port that has an access profile can access only those volumes that are defined in the profile. Depending on the capability of the particular host operating system, an access profile can contain up to 256 or up to 4096 volumes.

The setup of an access profile is transparent to the user when using the ESS Specialist to configure the host FC ports and ESS logical volumes in the ESS. Configuration actions that affect the access profile are:

- Configuring the host FC port in the ESS with its WWPN creates the access profile for that host FC port. Initially the profile is empty. That is, it contains no ESS logical volume assignments. The host FC port cannot access any ESS logical volumes that are already defined in the ESS.

- ▶ Adding new ESS logical volumes to the ESS, the new volumes go to the host FC port that is selected.
- ▶ Assigning ESS logical volumes to host FC ports, means adding them to the access profile of the selected host FC port and to pre-existing ESS logical volumes.

Removing a host FC port from the ESS, means deleting the host FC port definition and its access profile.

2.7.2 ESS managed ESS FC port WWPNS

Beginning with LIC level LBLD0312, there is a new function that allows the ESS to locally administer the WWPNS for the ESS FC ports. This means the ESS FC ports will have WWPNS which do not change when replacing an adapter. This is especially useful if SDD is used for multipathing and when zoning is implemented using the WWPNS of the ESS FC port.

Activating the WWPNS naming convention

The activation of this feature for all of the ESS adapters requires a 45 minute nonconcurrent service action and must be carried out by an IBM CE. Once this feature has been activated by the CE, this is done from the service menu by selecting:

1. Select **Configuration Options Menu**.
2. Select **Change/Show Control Switches**.
3. Fibre Channel WWPNS Method - Standard should be displayed, and if so, press the Enter key.
4. Press PF4, select **WWNN**, and then press Enter.
5. When instructed, IML the ESS.

From the time of activation onwards, all built in ESS FC ports and all new ESS FC ports will have the locally administered WWPNS. However, even if the feature is not activated, any replaced or added adapters will adopt the new locally administered WWPNS.

Host and switch considerations

After the CE has performed this service action and the naming convention has been successfully changed from Standard to WWNN, there are some actions that will need to be performed on the attached hosts and, in some cases, the fabric components, in order to allow the host FC adapter to successfully login using the new WWPNS.

These actions are host dependent. With respect to the fabric components, if WWPN zoning has been implemented, then it will be necessary to update every instance of the WWPN in the switches and/or directors.

The WWPNs of the adapter are built from the WWNN of the ESS and the location of the adapter in the ESS interface bays. This is done using the pattern shown in Figure 2-120.

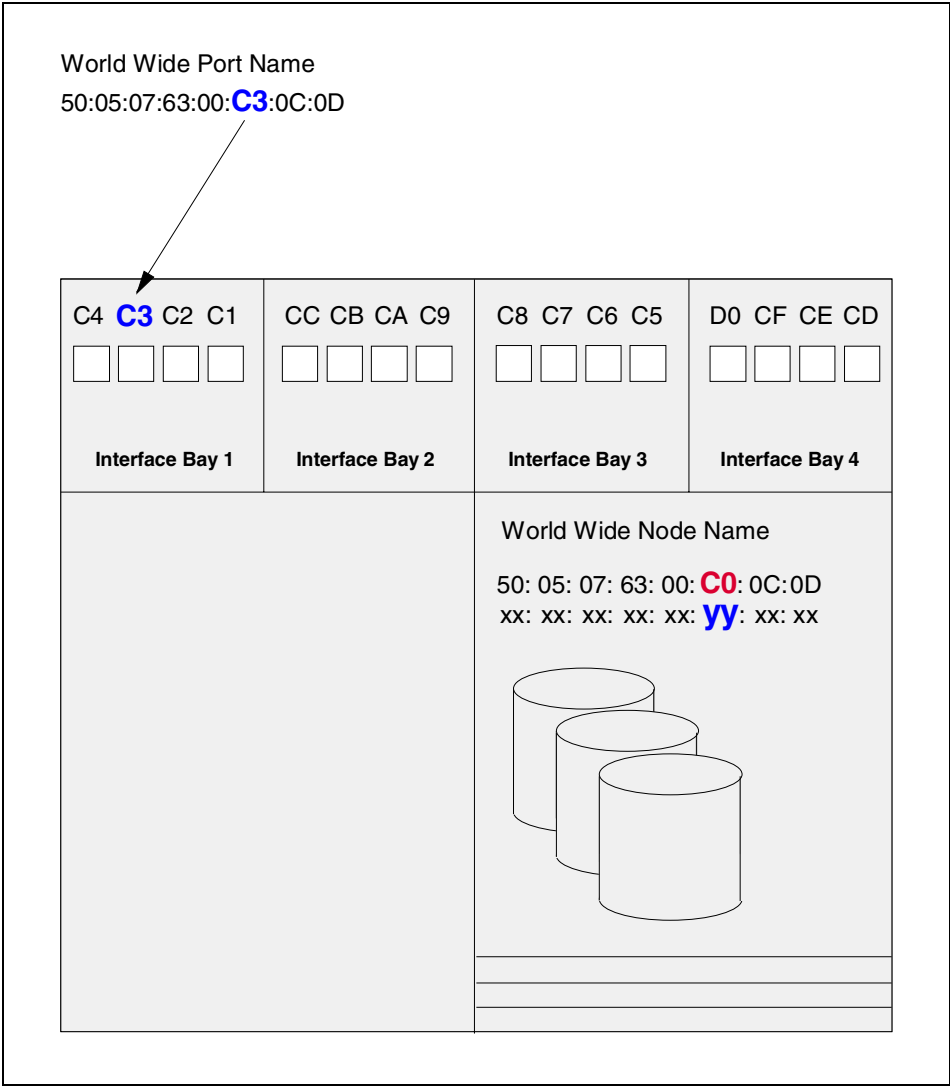


Figure 2-120 Pattern for the generation of the WWPNs of the ESS FC ports

The Information frame on the Storage Allocation — Graphical View panel reflects the new WWPNs of the ESS we have used for this illustration, which is shown in Figure 2-121.

Information	
Adapter Type: FC host attachment	
FCP Capable	
IO/Bay:	1
Adapter:	2
Port:	A
WWPN:	5005076300C30C0D
Topology: Point to Point	

Figure 2-121 Information Frame, reflecting the new locally administered WWPN

In Figure 2-122 we show the WWNN of the ESS that is used and modified when this feature is enabled.

 Welcome to IBM StorWatch Enterprise Storage Server Specialist	
Machine Type:	2105
Machine Model:	F20
Serial Number:	IBM-FCA13
WWNN:	5005076300C00C0D

Figure 2-122 WWNN of the ESS with locally administered WWPNs

What we also see in the information frame is that this is an ESS which has ESS FC ports that are also FICON capable. This adapter indicates that it is configured for FCP and not for FICON.

2.7.3 Configuring host FC ports for FC-SW and FC-AL

Now, as we have selected the first ESS FC port, we are going to configure this and two other ports to be able to communicate with our host FC ports. The first two ports we configure for use in a switched fabric environment and the second are for use in an arbitrated loop environment.

Selecting the topology for the ESS FC ports

The pull down list in the FC Port Attributes field as shown in Figure 2-123, is for selecting the topology the port is going to be connected to.




Figure 2-123 Pull down list, selected port currently undefined

The selected ESS FC port in bay one adapter three we start with is unconfigured. Because of this, we can choose by selecting between **Fibre Channel Point to Point** and **Fibre Channel Arbitrated Loop**. If a port is already configured for one topology and it has to be changed to another, the port must first be undefined. Then it is possible to choose the appropriate topology.

We can choose between **Fibre Channel Arbitrated Loop for FC-AL** and **Fibre Channel Point-to-Point for FC-PP** and for **FC-SW**. Some adapters — for example, the IBM FC6227 — operate in arbitrated loop mode even though they are directly connected to an ESS. This means that the ESS port will have to be configured in FC-AL mode.

As we have selected the topology for the selected ESS FC port, we see a small red line directly below the icon of the just-defined ESS FC port, as shown in Figure 2-124.

Configure Host Adapter Ports Reset Selected Port ?

Host Adapter Port: Bay 1, Adapter 3 Port A (FC)

Storage Server Attributes

Access_Restricted Fibre Channel Access Mode

FC Port Attributes

10000000C9251126 World-Wide Port-Name

Fibre Channel Point to Point Fibre Channel Topology

Figure 2-124 ESS FC port for Point to Point defined

This indicates that the port definition has been changed. We can now go on and select and define the other ports. The red lines will remind us where we made definition changes.

Now we go on and define port three in bay four for FC-SW and adapter three in bay two for FC-AL. The panel then indicates us what port definitions have been changed as shown in Figure 2-125.

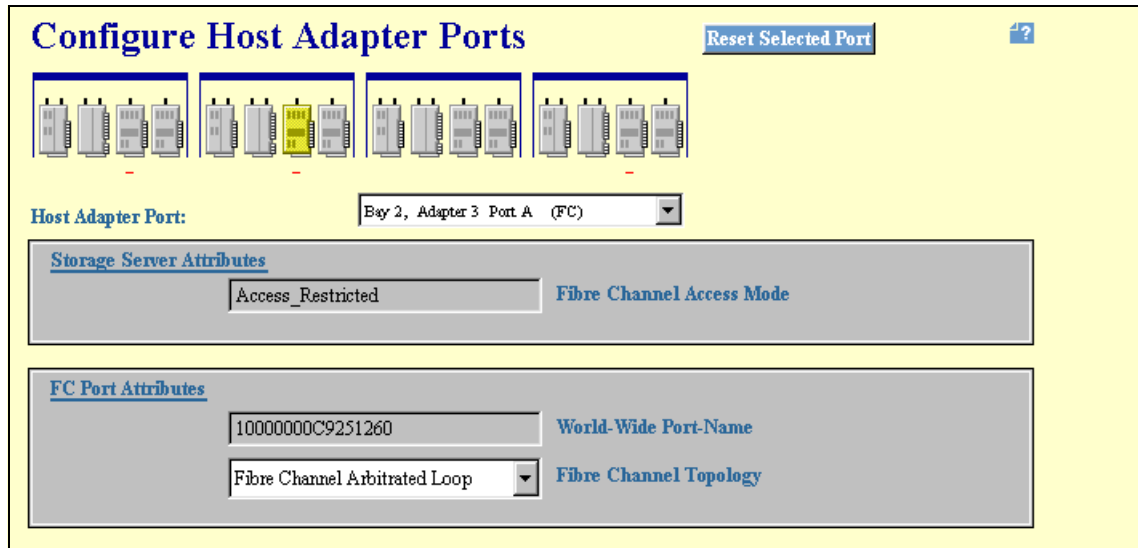


Figure 2-125 ESS FC port for Arbitrated Loop defined

Resetting changes made to a selected ESS FC port

Using the **Reset Selected Port** button on the upper right of the screen, we can reset the definition for the selected port to the status the port was in before we made any changes.

Reset Selected Port

This brings up a warning window as shown in Figure 2-126, stating that the configuration changes made are going to be discarded.




Figure 2-126 Warning window, configuration changes are going to be discarded

We do not want to reset our changes, so we are going to perform the changes.

Performing the configuration update

Click on the **Perform Configuration Update** button. During the process of performing the update we get a progress window as shown in Figure 2-127.

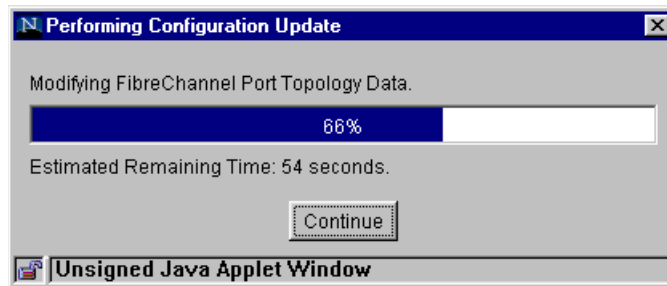


Figure 2-127 Progress window, changing the topology

Figure 2-128 shows a successful topology change of the ports.



Figure 2-128 Information window, successful topology change

Pressing the **OK** button takes us back to the Open System Storage panel.

Viewing changes in the Graphical View panel

The changed topology settings are reflected in the Storage Allocation — Graphical View panel. Now when we click on one of the ESS FC ports that we just configured, the information frame reflects the changes to the topology.

For port three in bay one, which is now configured for use with FC Point to Point, this is shown in Figure 2-129. For port three in bay two, which is now configured for use with FC Arbitrated Loop, this is shown in Figure 2-130.

Information	
Adapter Type:	FC host attachment
IO/Bay:	1
Adapter:	3
Port:	A
WWPN:	10000000C9251126
Topology:	Point to Point

Figure 2-129 Information frame, ESS FC port topology Point to Point

Information	
Adapter Type:	FC host attachment
IO/Bay:	2
Adapter:	3
Port:	A
WWPN:	10000000C9251260
Topology:	Arbitrated Loop

Figure 2-130 Information frame, ESS FC port topology Arbitrated Loop

We have now prepared the ESS to provide storage to our hosts.

We have arrived at the end of journey! Here is our last picture of the Storage Allocation — Graphical View panel, shown in Figure 2-131.

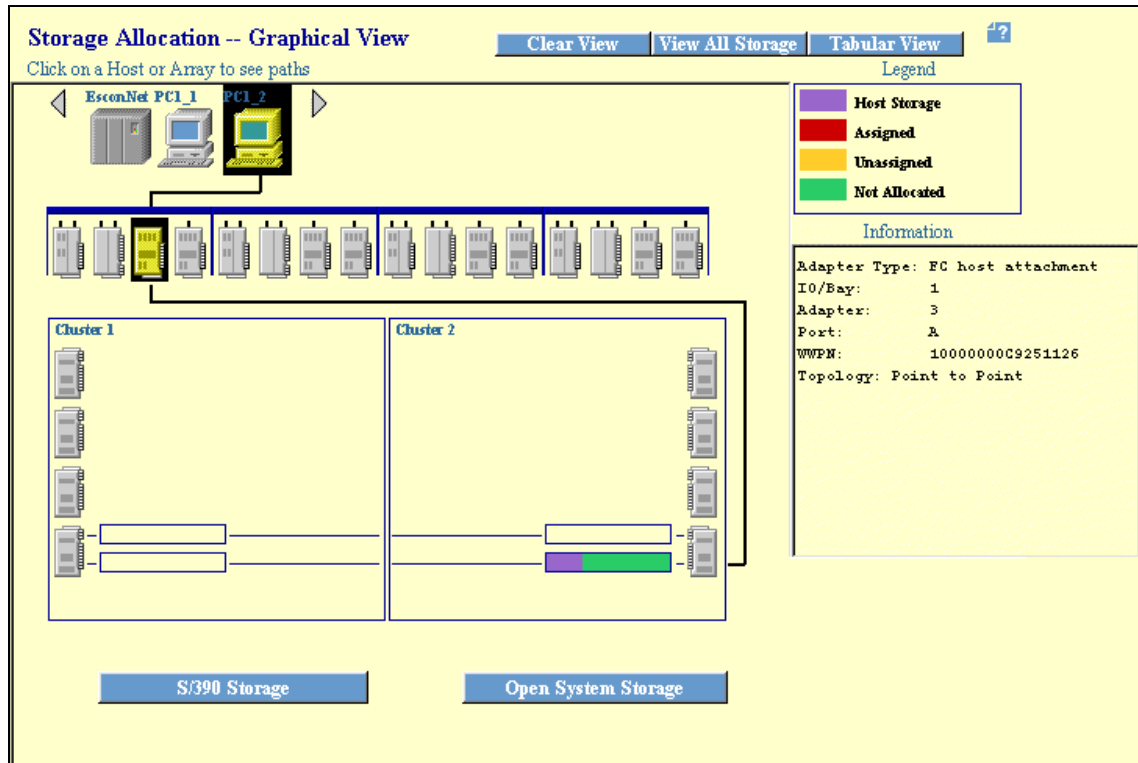


Figure 2-131 Storage allocation now completed

We see that the host FC port PC1_2 is restricted to see the one volume it has left through one FC port which is configured for point to point.

Related information: For more information on the ESS, see these references:

IBM Enterprise Storage Server Introduction and Planning Guide, 2105 Models E10, E20, F10 and F20, GC26-7294

- Introduces the product and lists the features you can order. It also provides guidelines on planning for the installation and configuration of the ESS.

IBM Enterprise Storage Server User's Guide, 2105 Models E10, E20, F10 and F20, SC26-7295

- Provides instructions for setting up and operating the ESS.

IBM Enterprise Storage Server Host Systems Attachment Guide, 2105 Models E10, E20, F10 and F20, SC26-7296

- Provides guidelines for attaching the ESS to your host system.

IBM Enterprise Storage Server SCSI Command Reference, 2105 Models E10, E20, F10 and F20, SC26-7297

- Describes the functions of the ESS. It provides reference information for UNIX and AS/400 hosts, such as channel commands, sense bytes, and error recovery procedures.

IBM Enterprise Storage Server System/390 Command Reference, 2105 Models E10, E20, F10 and F20, SC26-7298

- Describes the functions of the ESS and provides reference information for System/390 hosts, such as channel commands, sense bytes, and error recovery procedures.

ESS Web Interface User's Guide for ESS Specialist and ESS Copy Services, SC26-7346

- Provides instructions for using the IBM TotalStorage Enterprise Storage Server Web interface, ESS Specialist.

IBM Storage Solutions Safety Notices, GC26-7229

- Provides translations of the danger notices and caution notices that IBM uses in ESS publications.

IBM Enterprise Storage Server Configuration Planner, SC26-7353

- Provides work sheets for planning the logical configuration of the ESS. This book is not available in hard copy. This book is only available on the following Web site:

<http://www.storage.ibm.com/hardsoft/products/ess/refinfo.htm>

IBM Enterprise Storage Server Quick Configuration Guide, SC26-7354

- Provides flow charts for using the TotalStorage Enterprise Storage Server Specialist. The flow charts provide a high-level view of the tasks the IBM service support representative performs during initial logical configuration. You can also use the flow charts for tasks that you might perform during modification of the logical configuration. The hard copy of this booklet is a 9-inch by 4-inch fanfold.



Implementing the INRANGE FC/9000 Fibre Channel Director

The IBM machine type 2042 model number 001, INRANGE FC/9000 Fibre Channel Director, is the core product of a reseller agreement between IBM and INRANGE Technologies, which adds the INRANGE FC/9000 Fibre Channel Director to IBM's growing list of enterprise-class SAN fabric offerings.

3.1 Introduction to the INRANGE FC/9000 director

In this chapter we introduce the INRANGE FC/9000 Director and cover the following topics:

- ▶ Product architecture
- ▶ Zoning methodologies
- ▶ Management features
- ▶ Supported topologies
- ▶ Supported servers
- ▶ Supported devices

In Figure 3-1 we show the front view of an FC/9000.



Figure 3-1 front view of an INRANGE FC/9000 director

3.1.1 Product overview

The FC/9000 Fibre Channel Director is INRANGE's flagship SAN product. It was designed as a director-class product and has the following significant features:

- ▶ Support of Fibre Channel Protocol for open servers
- ▶ Support of FICON attachment for Mainframe class servers
- ▶ High port count, with up to 128 ports
- ▶ Fully redundant architecture with non-disruptive code load-capabilities
- ▶ Graphical user interface
- ▶ Broad range of supported port-topologies

Full loop support

INRANGE's Translative Loop Mode allows legacy nodes using loop protocols of any kind to be managed as if they were truly fabric nodes. This includes public and private loop ports with both initiators and targets.

We cover this in more detail in "Attaching loop ports" on page 220.

Port count

One chassis of an FC/9000 director can hold up to 64 GBICs. By interconnecting the main boards of two directors, a non-blocking 128 director can be created.

Due to the usage of GBICs instead of fixed port-cards, the granularity of port upgrades is one port.

Short wave and long wave ports can be mixed within all internal entities of this director.

Redundancy and non-disruptive activities

The FC/9000 features a N+1 redundancy for all of its elements. This redundancy results in high availability. Even in the case of failure of any internal elements, the director is still operational.

Obviously, in case of port-failures, additional multipathing software is needed to maintain this level of availability.

All activities, including code loads, code-activation, and replacement of failed parts can be done non-disruptively.

3.1.2 Currently supported zoning methodologies

Zoning is one of the key functions provided by SAN switches or directors. Basically it provides methods to limit possible communication between ports. This can be used to subdivide a physical SAN into separated logical ones.

Zoning can be seen as an important security feature of SAN products.

Since the definitions of different zoning-methods are not commonly agreed or adhered to within the SAN industry, we will refer to INRANGE's zoning terminology throughout this chapter. INRANGE offers two kinds of zoning:

- ▶ Hardware zoning
- ▶ Software zoning

Hardware zoning actually isolates ports from other ports within the same director. Consequently hardware zones are never overlapping. Any communication between hard zones is blocked. This is widely considered as the highest level of security.

Soft zoning allows you to overlap ports. That means it is possible to place ports into different zones.

Routing information is stored in a central name server table, which will be accessed by initiators before actually starting their I/O.

Soft zones within INRANGE reference to specific physical ports of the director. Soft zoning referencing to WWPN of attached nodes is planned to be made available by firmware upgrades in the second half of 2001.

You will find more detailed information on these topics in 3.4.6, "Defining name server zones" on page 252.

3.1.3 Management capabilities

FC/9000 fabrics can be managed using different communication methods:

- ▶ IN-VSN-Manager (IP based client/server management software)
- ▶ Simple network management protocol (SNMP)
- ▶ Serial interface (RS232, dedicated for INRANGE/ IBM engineers)
- ▶ Call Home (modem connection for notification purposes)
- ▶ Trivial file transfer protocol (TFTP) to load microcode IP-settings

Figure 3-2 shows an example of a director view with IN-VSN Manager.

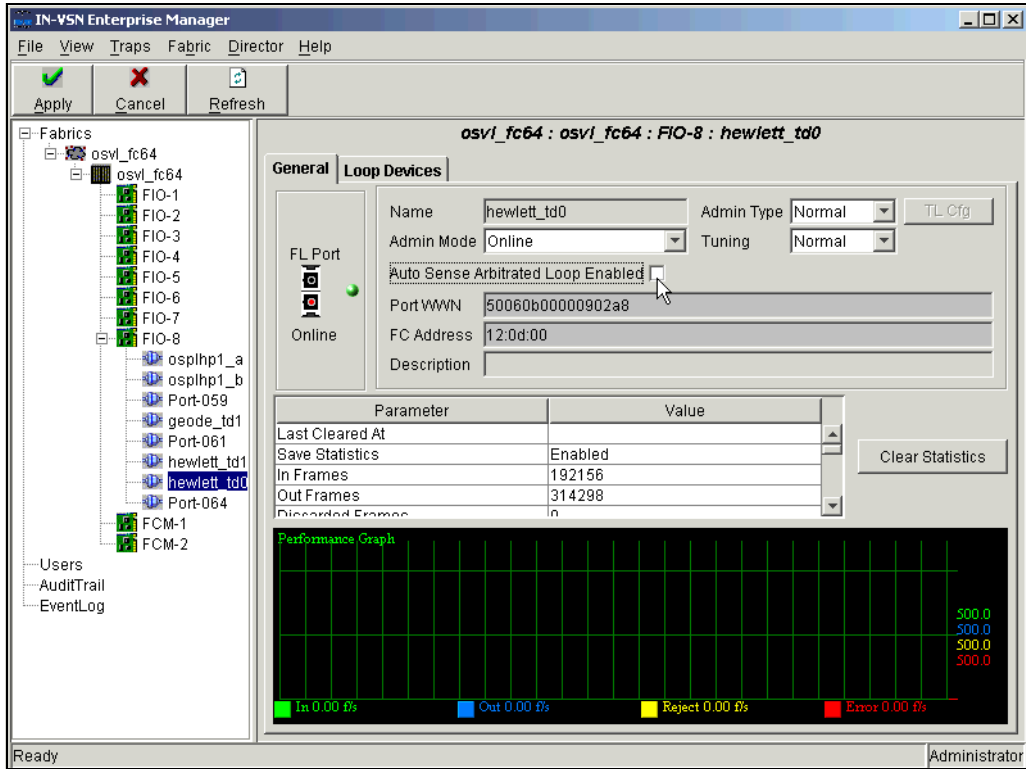


Figure 3-2 Example of IN-VSN management view

The most commonly used interface with the INRANGE FC/9000 director is the IN-VSN software tool. Therefore, in the following topics we will focus on how to use IN-VSN.

3.1.4 Supported protocols

In most environments, a homogeneous landscape of servers and storage is hard to find. For example, most tape-drives solely support FC_AL (known as arbitrated loop), whereas modern disk systems are widely used with point-to-point protocol (sometimes called P2P).

Thus, enterprise wide SAN solutions should be flexible enough to support different FCP topologies.

INRANGE does just that, by supporting these protocols:

- ▶ Open Systems:
 - Arbitrated Loop
 - Public loop
 - Private loop
 - FC-SW (Fibre Channel Switched Fabric)
- ▶ S/390 systems (zSeries):
 - FICON

All these can attached to a single INRANGE director at the same time.

Support of cascading

To create even larger fabrics FC/9000 directors can be cascaded. By doing this you can create fabrics with more than a thousand external ports. Today, IBM supports fabrics with up to 8 cascaded INRANGE FC/9000 directors.

However, cascading is not supported when using FICON attachments in the currently supported microcode releases.

3.1.5 Supported host attachment

Currently these servers are supported by IBM to attach to the FC/9000:

- ▶ IBM 9672 Enterprise G5, G6 and z-series server with FICON Channel Card (FC2314 and FC2316), Longwave and Shortwave laser
- ▶ IBM Netfinity server with Windows NT 4.0 (SP6a) and FASTT Host Adapter (P/N:00N6881)
- ▶ Intel-based servers with Windows NT 4.0(SP6a) or 2000K and QLogic QLA2200F Host Bus Adapter
- ▶ Sun Servers with Solaris 2.6,2.7, 8 and Host Bust Adapters - SBus JNI FC64-1063-N or PCI LP8000, QLA2200F, JNI FCI-1063

Additionally, the following servers are in their testing process now, and are expected to be supported by IBM once testing is complete:

- ▶ IBM RS/6000 and RS/6000 SP Servers with AIX 4.3.3 and Gigabit Fibre Channel Adapter (FC 6227)
- ▶ HP/9000 Servers with HP/UX 11.0 and A5158A Adapter

For the latest support matrix of INRANGE FC/9000, refer to:

http://www.storage.ibm.com/ibmsan/products/directors/prod_data/supserver-042.html

3.1.6 Supported device attachment

Currently these Storage products are supported by IBM:

- ▶ IBM Magstar 3590 A60 Tape Subsystem with FC 3432 (FICON Attachment - Longwave) in a 3494 Tape Library, in a 3590 Model A14 Frame, in a standalone 19-inch rack; or C14 Silo Compatible Frame
- ▶ IBM Magstar 3590 Tape Subsystem 3590-E11/B11 in standalone rack mount; 3590-E1A/B1A in 3494 Tape Library or 3590C-Silo Compatible with FCP-Attachment
- ▶ IBM Enterprise Storage Server (M/T: 2105-Fx0) and FICON Host Adapter
- ▶ IBM Enterprise Storage Server (M/T: 2105-Ex0 or 2105-Fx0) and Fibre Channel Host Adapter (FC 3022)

Other attachable devices include HDS 7700E, HDS 99x, EMC SYM 5x and Sun A5200. However, these devices are directly supported by INRANGE as a part of the IBM-INRANGE cooperation.

For the latest INRANGE supported device list, refer to:

<http://www.inrange.com/ibm/interoperability.php>

For the latest IBM supported device list, refer to:

http://www.storage.ibm.com/ibmsan/products/directors/prod_data/supserver-042.html

3.2 Installing an INRANGE SAN

The following chapter covers how to install an INRANGE SAN including topics like:

- ▶ Target environment
- ▶ Prerequisites
- ▶ Initial setup
- ▶ Installing the IN-VSN management software
- ▶ Accessing the IN-VSN management software
- ▶ User management

3.2.1 Target SAN environment

In our target ITSO environment we used these components:

- ▶ One Enterprise Storage Server 2105 F20
- ▶ One INRANGE Fibre Channel director FC/9000, equipped with 64 GBICs
- ▶ Four Hewlett Packard servers running HP-UX 11
- ▶ ShortWave fibre infrastructure with 50 micron core diameter

In Figure 3-3 we show our target ITSO environment.

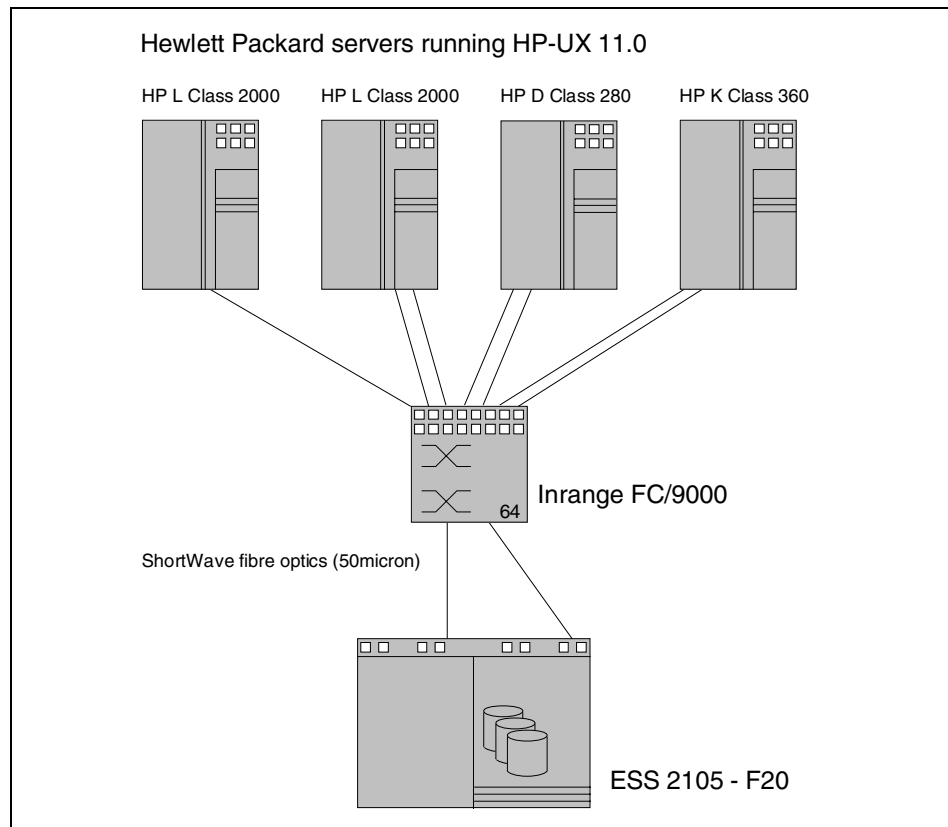


Figure 3-3 Initial INRANGE environment

For management purposes we have one dedicated PC server attached to the director. This management station runs the server portion of the INRANGE IN-VSN Manager to control all fabrics.

3.2.2 Verifying all prerequisites

Before starting the initial setup, it is important to have performed these steps:

- ▶ Obtain two power supplies for each director, for redundancy reasons.
- ▶ Ensure that you have the capability to attach directors and management station to ethernet network.
- ▶ Prepare a stable floor ground — one director weighs from 8 to 95 kg.
- ▶ Use the official INRANGE 19" rack, or any 19" racks that meet NSI/EIA 230-standard.

3.2.3 Initial setup of INRANGE FC/9000 IP settings

The INRANGE Director is delivered with the current supported level of firmware and a default TCP/IP setting of 10.1.1.51 and 10.1.1.52, and a Subnet mask setting of 255.255.255.0. There is also a default chassis ID, switch ID and fabric ID. Without doing any kind of cascading, there is no need to change these IDs.

It is the responsibility of an IBM Customer Engineer (CE) to reset all default addresses to reflect the environment in which it is being installed.

For added security, the TCP/IP address can only be set or reset by the CE, using an RS232 connection and entering the CE userid and password. The new TCP/IP address will be displayed in an LED panel that can be located on the FCM module.

The PC that will be used for the IN-VSN server needs to be connected to the FCM modules using the Ethernet ports.

You should also ascertain whether the applicable microcode level is installed or needs to be installed. The same is true for the IN-VSN management software:

- ▶ The director microcode level that we will be using is: 02.00.
- ▶ The IN-VSN management tool release is: 2.0.0 (Build RC6).

Since these codes are subject to change as new functions are added and improvements made, ask your IBM or INRANGE Technical contact what the current supported levels are.

3.2.4 Establishing network connection

As mentioned above, the INRANGE director and its management PC are delivered with pre-installed IP settings. They are efficient enough to set up a small private network with just directors and the management PC as members.

You can leave the initial IP setup as it was delivered for use as, or like, a private network. Consequently only local users can attach to the directors and its IN-VSN management tool.

Or you can adopt your corporate network settings to enable remote IN-VSN access.

Leaving all IP settings as a private network

For using the IN-VSN software from this local private network, the supplied 3Com/US Robotics hub is sufficient enough. To enable all IP based components to talk together, plug in the Management PC and both INRANGE IP ports to this 3Com Hub. So initially, our network setup looks like that shown in Figure 3-4.




Figure 3-4 Private IP network for initial IN-VSN management ability

Enabling IN-VSN access from a corporate network

To exploit the remote management capabilities of the IN-VSN management software, we recommend that you connect this network as a subnet to the corporate LAN.

Before actually changing any parameters, it is important to obtain all the information you will need in advance, such as:

- ▶ Available IP addresses
- ▶ Valid subnet mask
- ▶ Default gateway

To connect the INRANGE subnet to your corporate LAN, follow these steps:

1. Change the IP setting of the primary FCM blade using the RS232 Interface (this should be done by INRANGE or IBM Customer Engineers).
2. Change the IP setting of the secondary FCM blade using the RS232 interface (this should be done by INRANGE or IBM Customer Engineers).
3. Change the IP settings of the management PC using the common Windows Control Panel tools. Reboot the Windows server.
4. From the management PC, ping both IP addresses of the director to ensure that everything is set properly (provided that all components are still connected to the 3Com IP-Hub).
5. Then attach all required ports to a switch or hub that is actually connected to the corporate network. You could still use the 3Com Hub as a connecting device.

After attaching our INRANGE setup to IBM's corporate network, as shown in Figure 3-5, we are now able to access the IN-VSN software from wherever we are in the corporate network.




Figure 3-5 INRANGE setup attached to a corporate network

To gain actual access to the IN-VSN management tool, you will need to install an IN-VSN client that can then communicate with the IN-VSN server over the corporate network.

Instructions on how to install both the server and the client part of IN-VSN software are provided in 3.2.5, “Installing the IN-VSN Enterprise Manager software” on page 187.

Setting up high security network access

In the previous setup we described a network layout in which the Management PC as well as the fabric components were connected to the corporate network.

However if this corporate network itself cannot be considered as secure enough, we recommend that you separate the fabric components from the corporate network. The only bridge between such a separated fabric management network and the corporate network would be the Management PC.

Referring to Figure 3-6, we see that now only the Management PC can access the directors's IP ports.

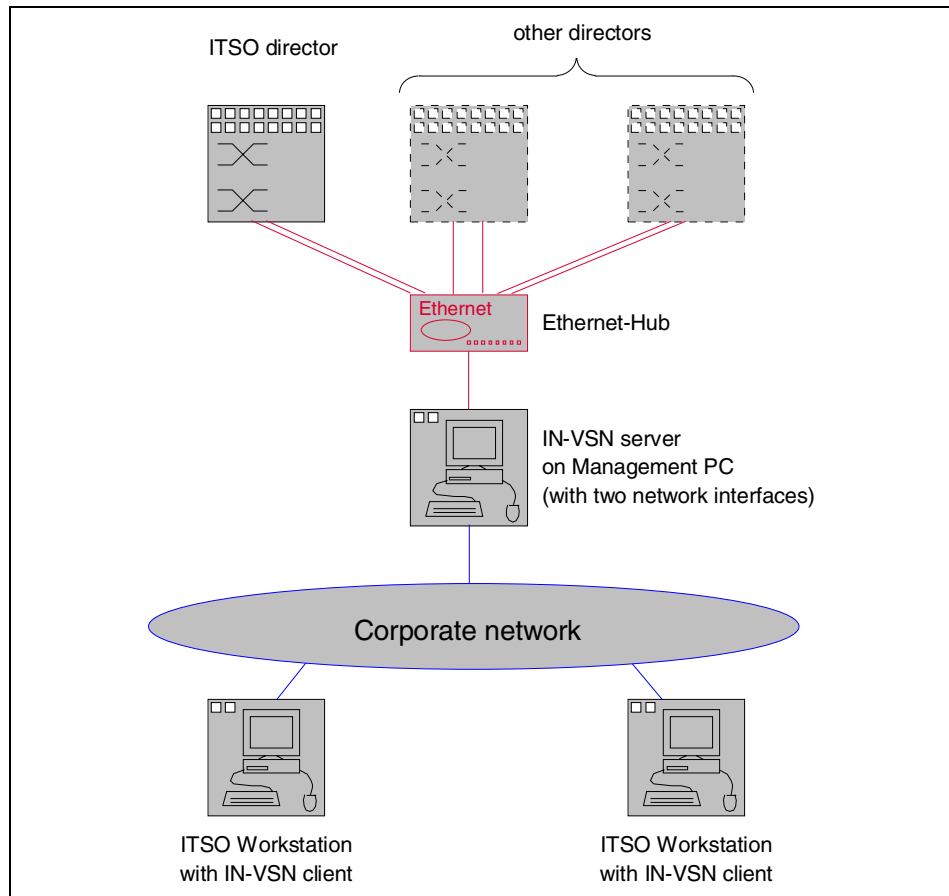


Figure 3-6 INRANGE setup with secure director access

Direct IP access from the corporate network to directors is now impossible. The only way to gain access is using the IN-VSN server.

We consider this as the most secure network setup for remote IN-VSN access. However, you will need two network interfaces in the Management PC.

3.2.5 Installing the IN-VSN Enterprise Manager software

There are a number of different ways in which the IN-VSN software can be delivered.

If software feature code 7600 is configured, the workstation pack will include a PC, monitor, Ethernet Hub, and Modems. The management software comes pre-installed on the PC.

If you require more than one version of the IN-VSN client code, this is achieved by using feature code 7201.

If no additional software feature codes are configured when ordering the INRANGE Director, you will receive a CD that contains the IN-VSN server and client code. This CD is only licensed for one copy of the server and one copy of the client.

PC hardware for IN-VSN server

The server and client can be installed on a customer owned and designated PC, as long as it meets with the following minimum recommended specifications:

IBM Pentium III PC or an IBM-Compatible Pentium III PC

- ▶ 500 MHz processor
- ▶ 128 MB SDRAM
- ▶ 4 MB Video RAM
- ▶ 13 GB hard drive
- ▶ 1.44 MB diskette drive
- ▶ CD-ROM
- ▶ 1 parallel port
- ▶ 1 Ethernet 10Base-T/100 Base-TX
- ▶ IBM or equivalent mouse
- ▶ 2 external serial com ports (debug serial cable, external modem)
- ▶ 1 internal com port (internal modem)
- ▶ 17-inch 1280 x 1024 x 256 SVGA monitor
- ▶ Windows NT 4.0 with Service Pack 6a and higher
- ▶ Windows 2000 Professional
- ▶ Internal and External modem compatibility
- ▶ Analog protocol support: V.90 5 6kb/s ITU Standard, 3Com/US Robotic/Multi-Tech (international) hardware compatible modems recommended (do not use Win modem software controlled modems)

Be aware that only Windows operating systems are supported as software platforms for IN-VSN tools.

Running the setup procedures of IN-VSN server and client

Both the IN-VSN server and client require the following two programs to be started from the installation CD:

- ▶ Use **j2re1_3_0-win.exe** to install the appropriate Java programs.
- ▶ Use **setup.exe** to actually install the IN-VSN server and/or the IN-VSN client software.

For systems on which the IN-VSN server should run, an additional **odbc** driver has to be installed:

- ▶ Use **odbcinst.exe** to accomplish this.

To run the IN-VSN installation, follow these steps:

- ▶ From a Windows browser or a command prompt, start **D:\setup.exe**.
- ▶ The install wizard will appear. Click **Next** to continue the installation as shown in Figure 3-7.



Figure 3-7 Running the IN-VSN setup: Initial screen

The setup program then checks whether you have the appropriate level of Java installed on this system or not, as shown in Figure 3-8.

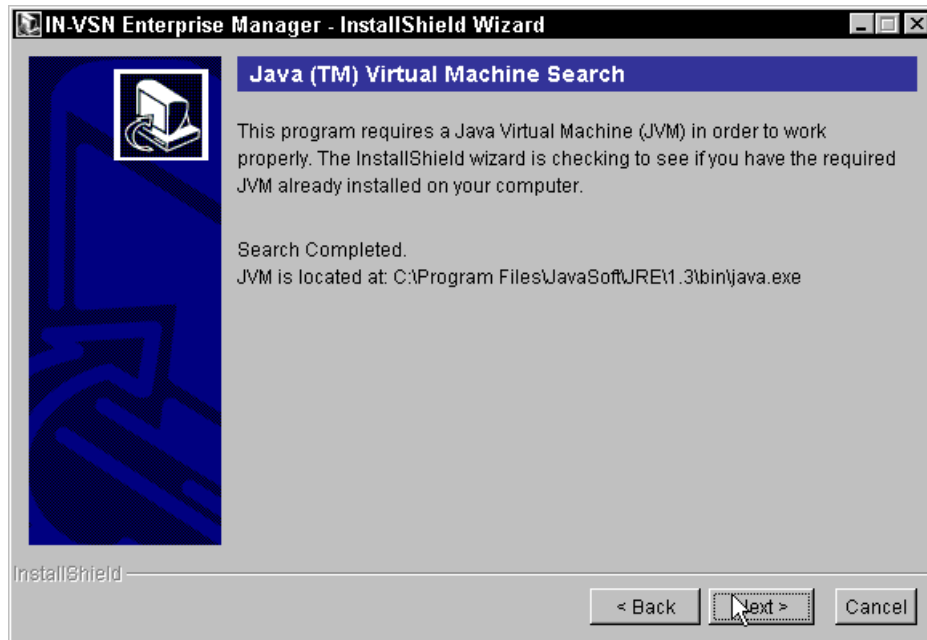


Figure 3-8 Running the IN-VSN setup: Java VM search

Click **I Agree** and **Next** in the Licence Agreement Window.

Then you must decide which part of the IN-VSN software should be installed, server and/or client, as shown in Figure 3-9.




Figure 3-9 Running the IN-VSN setup: Feature selection

After selecting the features that you wish to install, you may either accept the default installation path or enter a new one. The default installation path for both client and server is C:\INRANGE\Fc9000. The setup program automatically creates the folder if you agree to its creation.

Before actually starting the installation a list of all installation parts is shown in a list with its needed disk space. Click **Install Now** to accept these settings, as shown in Figure 3-10.



Figure 3-10 Running the IN-VSN setup: Verifying all packages to install

After installing these IN-VSN packages, you are given an installation summary which tells you that all parts were installed successfully.

Running the setup procedure to install the odbc driver

As mentioned before, all systems that should run the IN-VSN server part need to have an appropriate **odbc** driver. For IN-VSN clients this part is not needed.

To install it, run **odbcinst.exe**, which is part of the IN-VSN installation CD.

Leave all components selected in the **odbc** installation window, as shown in Figure 3-11.

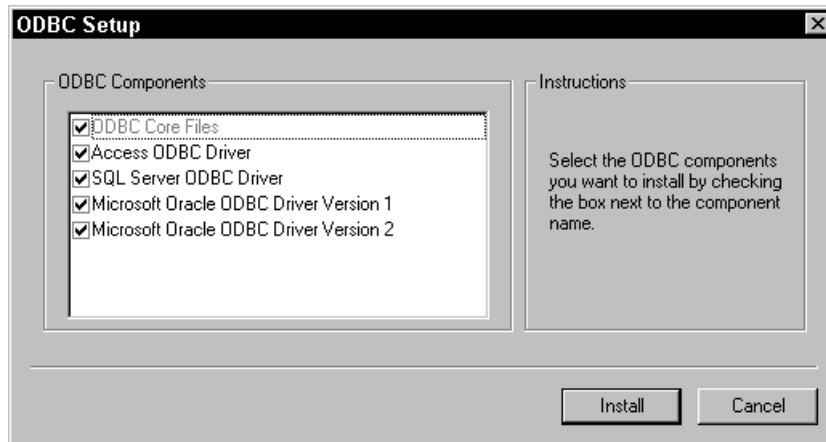


Figure 3-11 Running *odbc*

3.2.6 Accessing the management tool

Make sure that the IN-VSN server is running either on the local machine or elsewhere in the network.

Since the IN-VSN management tool is not browser based, it is necessary to have the client piece of software installed on your workstation.

We installed the IN-VSN client software on two different workstations in our target environment, as shown in Figure 3-12.

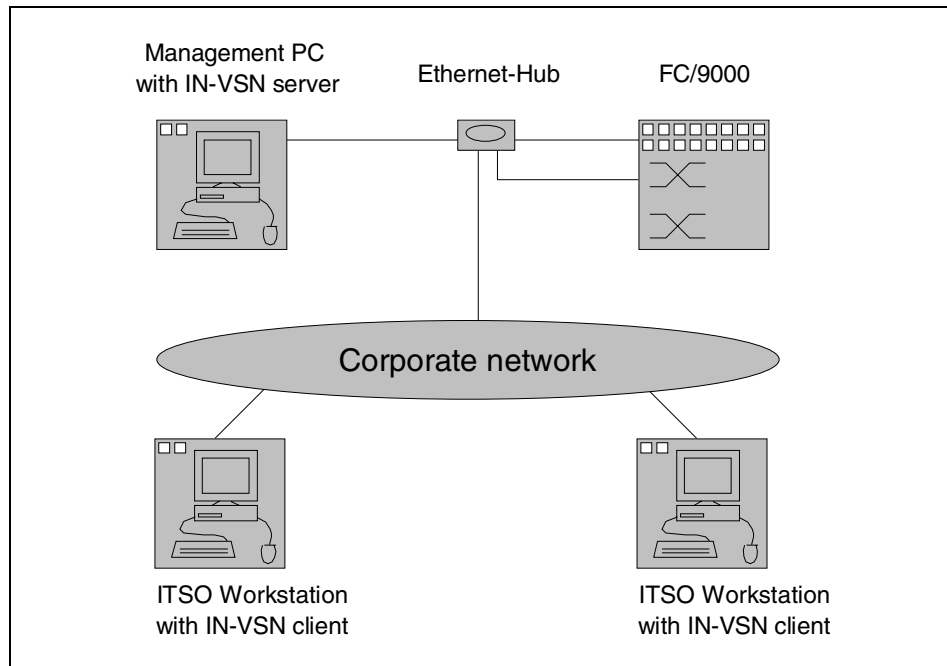


Figure 3-12 Final IN-VSN software setup

This layout makes it possible to have multiple users logged in to the IN-VSN software at the same time.

IN-VSN clients do not talk directly to the FC/9000. All communication to the FC/9000 is done via the IN-VSN server. However, the user interface is the IN-VSN client.

All director or fabric management related tasks are done using the IN-VSN client interface. The server interface is only used for initial database backup and Call Home settings.

Starting the IN-VSN server

Before you can actually start any IN-VSN client sessions, the IN-VSN server must be up and running.

You can start the IN-VSN server on the management PC by double-clicking the **IN-VSN Enterprise Mgr** icon on the desktop or using the Windows Start Menu.

For starting the IN-VSN server, no user login or password is required.

Once the IN-VSN server is up and running, an IN-VSN server window should appear with server running status, as shown in Figure 3-13.

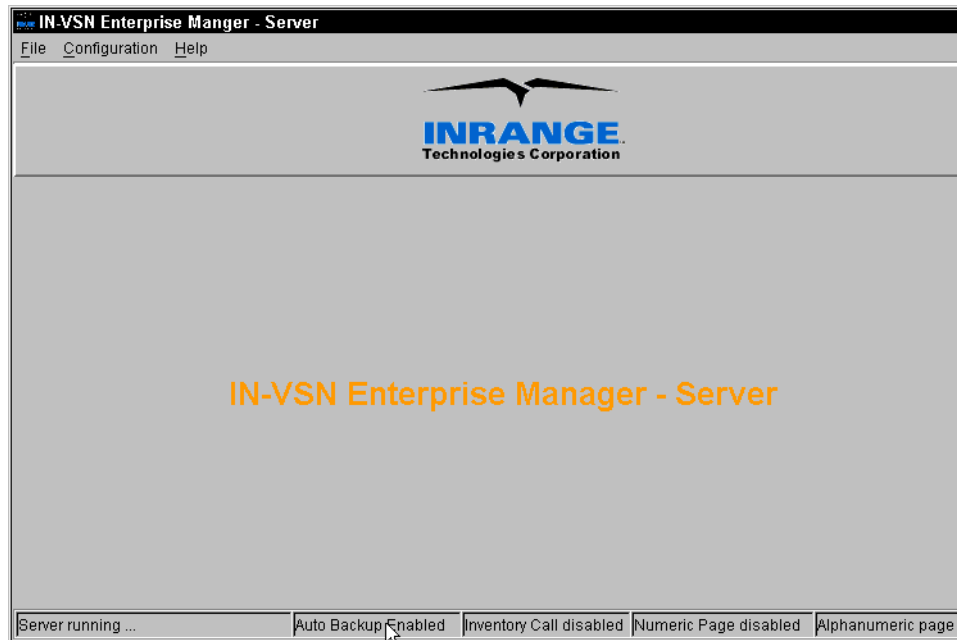


Figure 3-13 IN-VSN server up and running

Another indicator of the running status is the permanent change of colors of the text in the middle of the IN-VSN server window. Once the IN-VSN server is running, you can then start multiple IN-VSN client sessions pointing to the servers IP address or DNS name.

Starting IN-VSN client sessions

The following prerequisites must be fulfilled before being able to start an IN-VSN client session:

- ▶ Have the IN-VSN client code installed.
- ▶ On the workstation running the IN-VSN client, have the appropriate Java Virtual Machine installed.
- ▶ Have an IN-VSN server running locally or somewhere else in the network
- ▶ Have IP connection from IN-VSN client to IN-VSN server.

To start the IN-VSN client, just double-click the **IN-VSN client** icon on your desktop or use the Windows Start menu.

After starting the client, a login window will appear. Enter user name (default *admin*) and password (default *admin*) as shown in Figure 3-14.

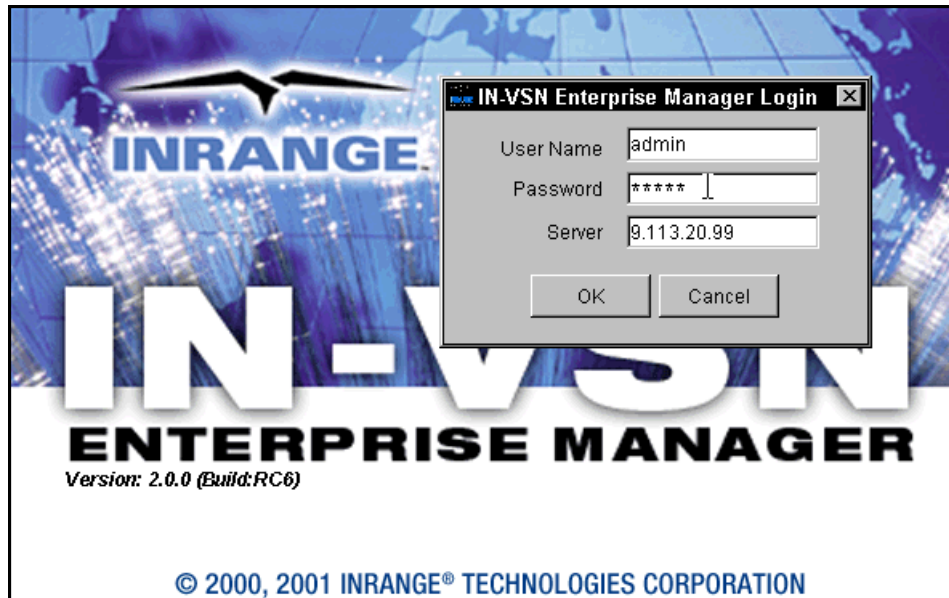


Figure 3-14 Starting an IN-VSN client session

Additionally, you will have to enter the IN-VSN server destination. This is where the client software should find the IN-VSN server running. If the server is running on a different machine, then enter the IP address of the IN-VSN server. If both pieces of software run on the same machine, just select **localhost**.

When logged into the IN-VSN server as the first client ever used, then all settings are still default values and no specific fabric or director information is available at all. This is shown in Figure 3-15.

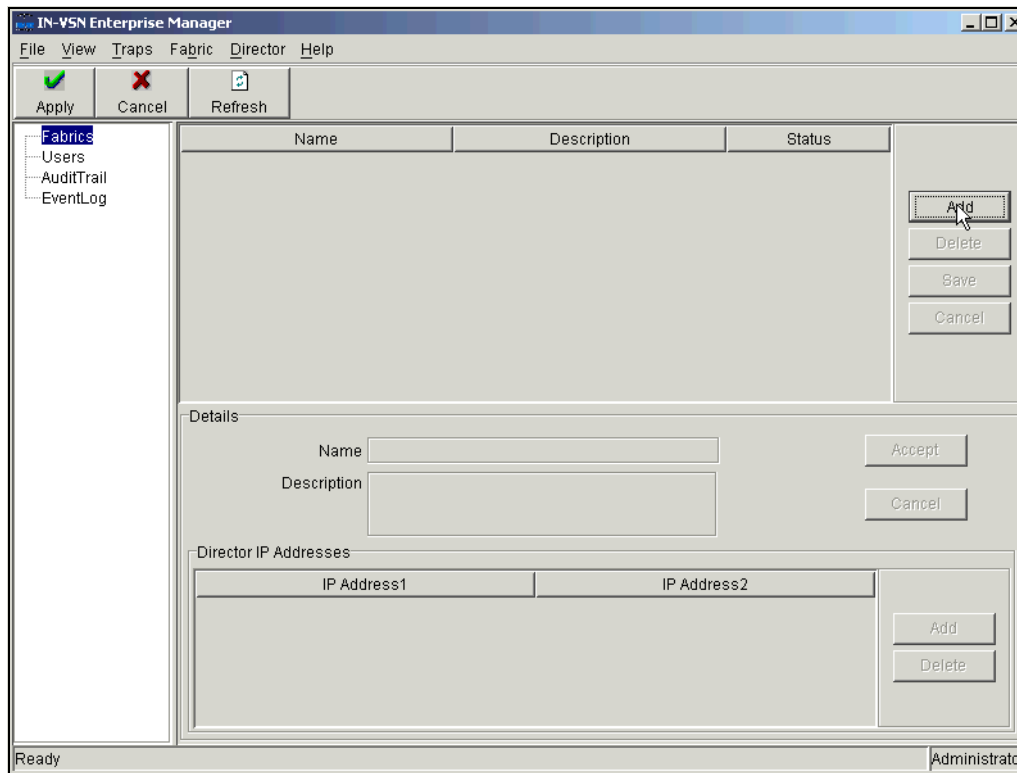


Figure 3-15 Initial screen of an IN-VSN client session

Once you are logged in to an IN-VSN client session, you can then customize and set up your fabric layout. These tasks include:

- ▶ Definition of IN-VSN users
- ▶ Setting up operational parameters
- ▶ Defining fabrics
- ▶ Mapping of directors
- ▶ Naming of directors and ports
- ▶ Creation of zones

These tasks are described in the following topics.

3.2.7 Defining users

Different user-levels are provided by INRANGE. You can have multiple users with different levels defined at a time. Default user names are the same as their level, and are as shown in Table 3-1.

Table 3-1 User levels and default users

default user	default password	explanation
viewer	viewer	view and monitoring
oper	oper	operations like hard zoning, naming or port definitions
admin	admin	any zoning, naming, definitions, user-management
maint	won't tell You	special user for customer engineer (CE) access

These default users and passwords are, like other products, the same for all IN-VSN servers. We strongly advise you to change them as soon as you start with your fabric definition to avoid any unauthorized and unwanted access.

Adding users

To enter the user definition screen, click the **Users** tab as shown in Figure 3-16.

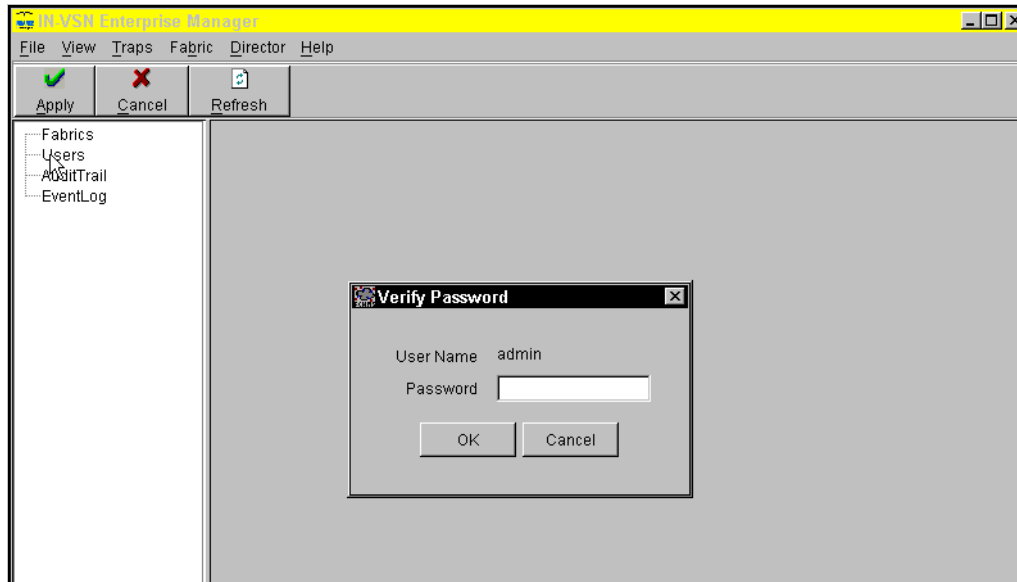


Figure 3-16 Adding an IN-VSN user: Re-enter password

For security reasons, you will have to re-enter your password.

After re-entering the password, the current user definitions are displayed. Initially, three users should appear as already saved:

- ▶ Admin
- ▶ Oper
- ▶ Viewer

Click **Add** to enter the information for a new user as shown in Figure 3-17.

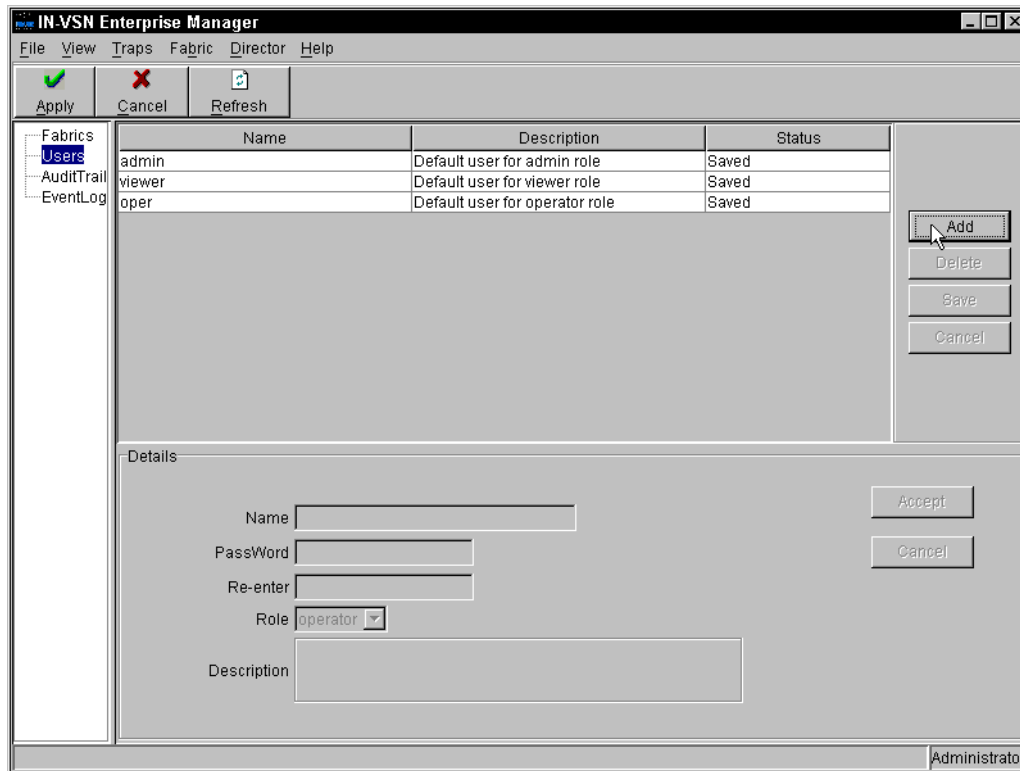


Figure 3-17 Adding an IN-VSN user: Overview

After clicking **Add**, the Details fields become active and you are able to enter the data needed to create a new user. You can give all rights to this new user apart from Maintenance level (maint).

In Figure 3-18 we create a new user called itso_1 with admin rights.

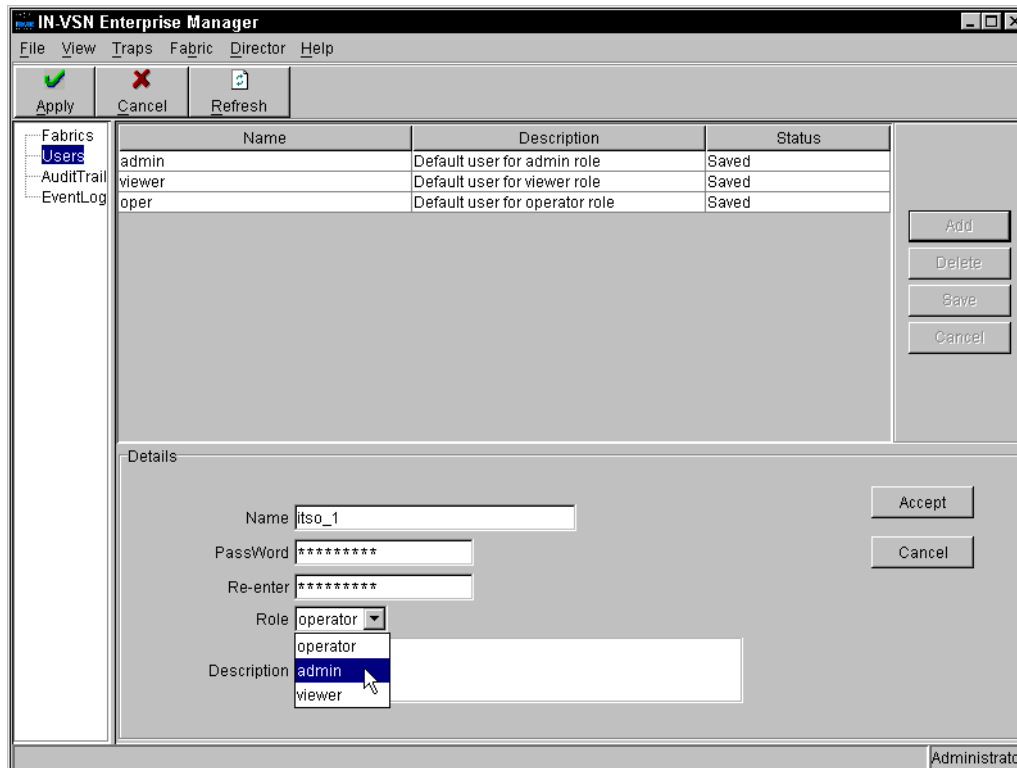


Figure 3-18 Adding an IN-VSN user: Creating an admin user

Click **Accept** to add this newly created user into the list of already known users.

However, this new user is not saved until you click **Save**, as shown in Figure 3-19.

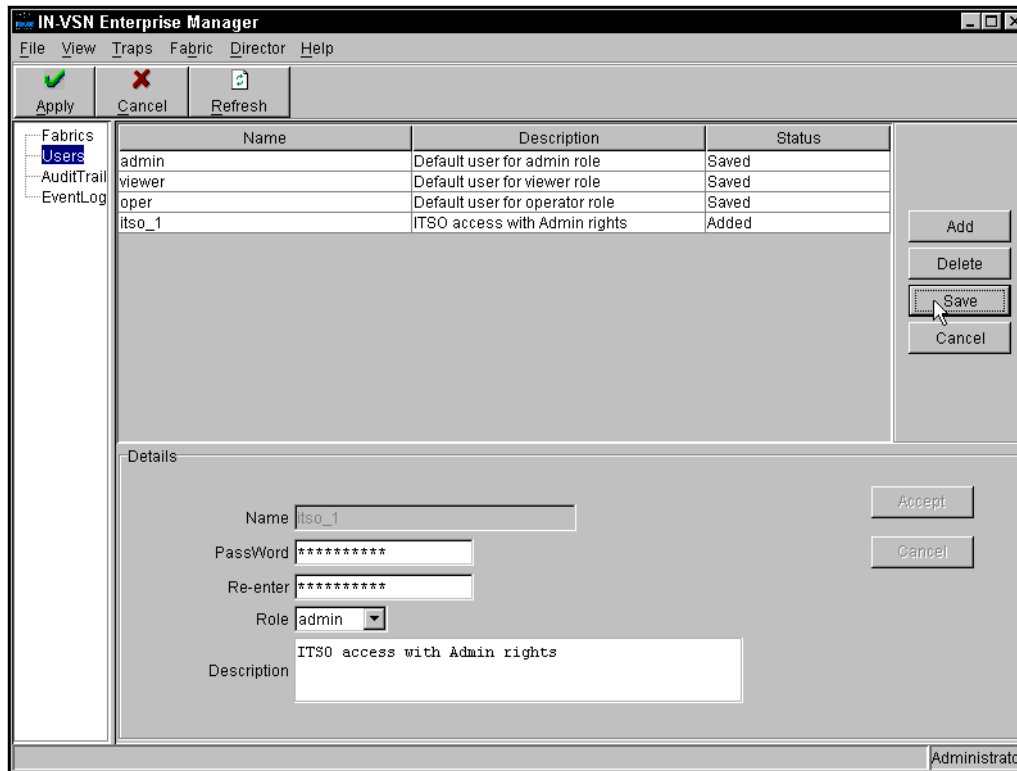


Figure 3-19 Adding an IN-VSN user: Saving new user configuration

You must take into account that the default users will not be deleted after adding new users. This may be different to other products that you are familiar with. To delete the default users, you can use the same procedure as for all other users. How to delete users is described in the topic that follows.

Deleting users

Deleting users is done using the same panel as we used for adding users. Note that you cannot delete the user that you are logged in with.

To remove any user, select the victim in the list of already defined and saved users. Click **Delete** and confirm this action as described in Figure 3-20.

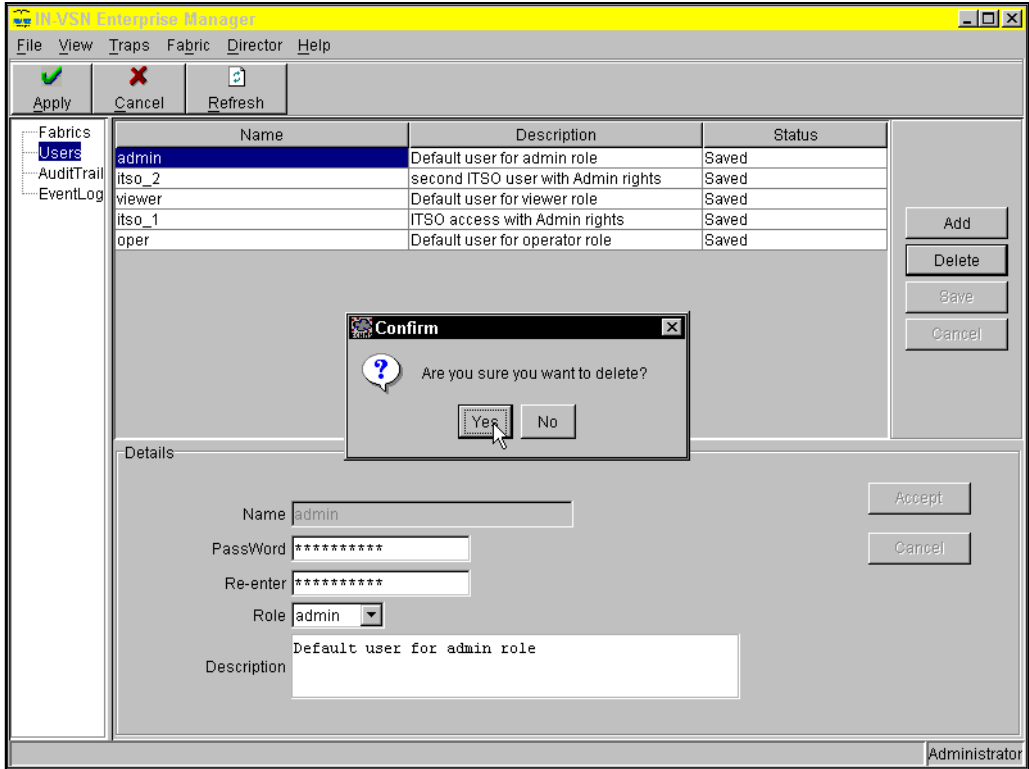


Figure 3-20 Deleting an IN-VSN user: Confirmation

Once you have confirmed this deletion, that particular user is deleted in the IN-VSN database of the IN-VSN server. It is also removed immediately from the user list of the IN-VSN client as shown in Figure 3-21.

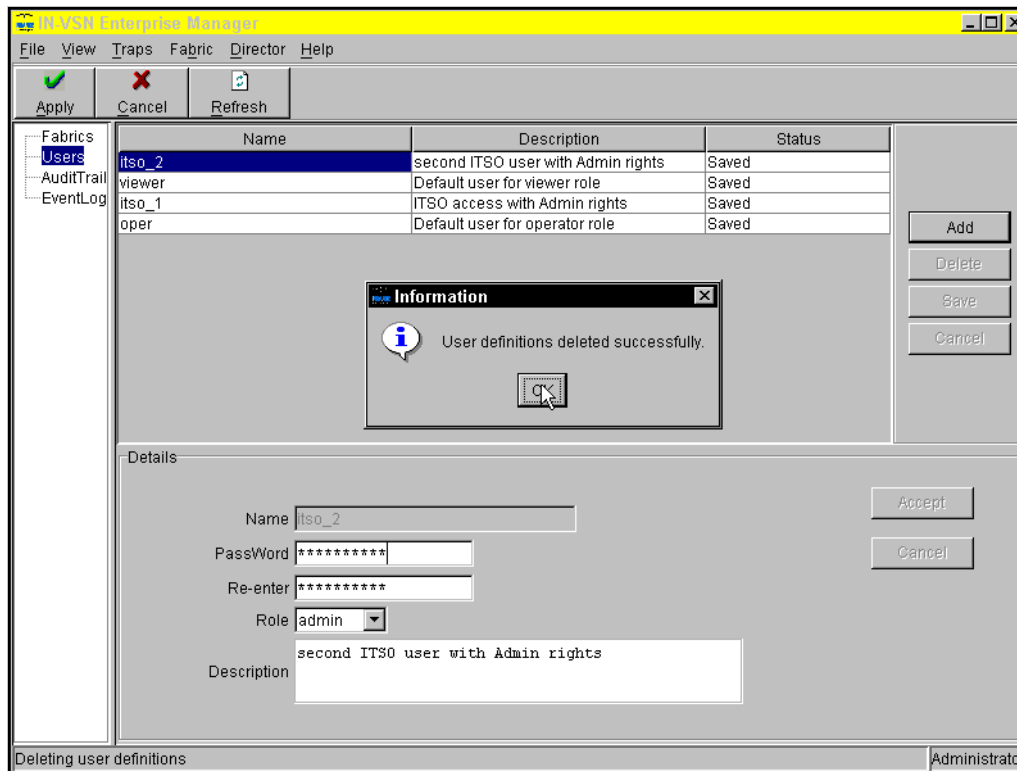


Figure 3-21 Deleting an IN-VSN user: Updated user list

Changing user definitions

Once you have added users, you can change their attributes. All attributes, excepting the name, can be changed. These are:

- ▶ Password
- ▶ User rights
- ▶ Description

To change a user, first select it in the **Users** list. Once you have selected a user, you are automatically in change mode. Change the desired settings and click **Accept** as shown in Figure 3-22.

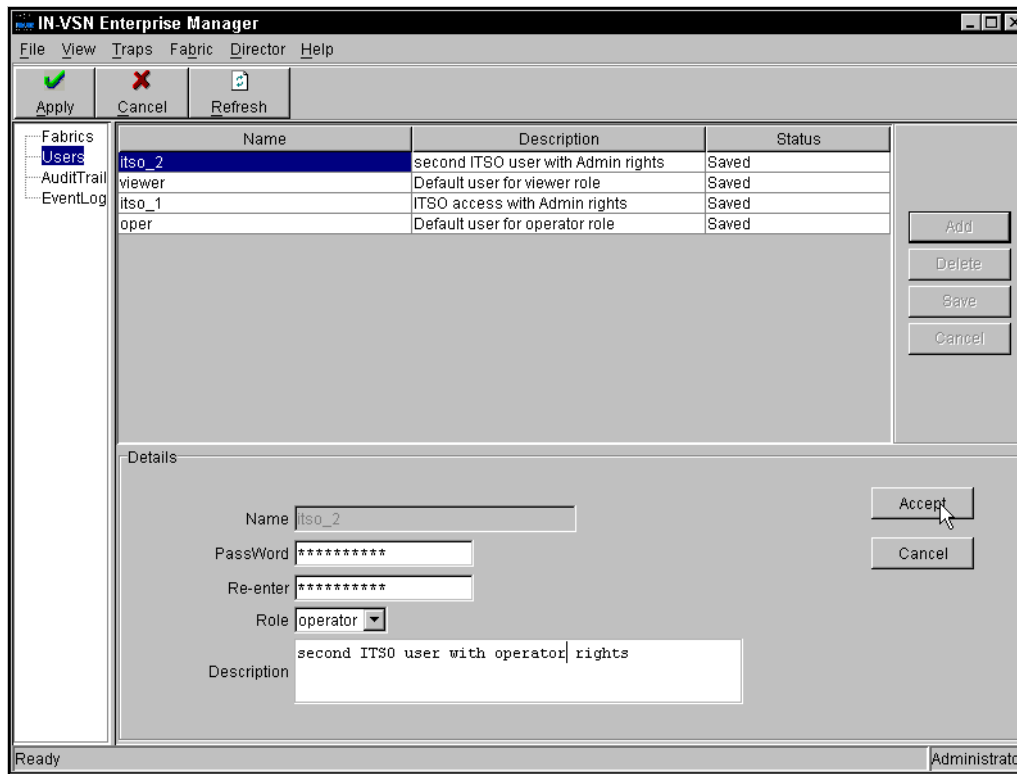


Figure 3-22 Change attributes of an existing IN-VSN user

We have changed the rights of user itso_2 from admin to operator.

After accepting the changes, this modified user will appear with updated status. You can modify multiple users before actually making these changes effective by clicking **Save**.

To be able to change all users, you must be logged in with admin rights. If you are logged in with only operator rights, then only your own user's password and description can be modified. An operator is not allowed to change its own rights (operator to admin, for example).

3.2.8 Security considerations

As a SAN may store a great deal of sensitive and often confidential corporate information, it might be interesting to consider some security issues, such as:

- ▶ User management
- ▶ Consequences of zoning methods
- ▶ Location access

We will describe these issues as they relate to the usage of the INRANGE FC/9000 director.

User management

As described in 3.2.7, “Defining users” on page 197 INRANGE provides different types of users.

User settings are saved in the IN-VSN servers database but not on the directors themselves. To ensure the ability to recover from an IN-VSN server failure, we advise you to use the Auto-Backup function. This backup setting is done using the IN-VSN server interface as described in 3.3.1, “Setting up operational parameters” on page 207.

You can have multiple users defined and also logged into the IN-VSN software at the same time. Please be aware of these important issues:

- ▶ You cannot segment the fabric for different users. So an IN-VSN user always has access to the whole fabric. It is not possible to limit the user’s access to NT server ports or AIX ports only. The philosophy behind this is the same as with Brocade or McDATA’s management tools. A SAN should be managed by a dedicated storage manager, regardless of which server platforms are attached to the SAN.
- ▶ Different users with the same level (for example, admin) can have access to the fabric at the same time. This could result in chaos, since both active users are allowed to make changes. Consequently, we recommend that only a limited number of people have admin rights to the same fabric.
- ▶ The IN-VSN default users (admin, oper, view) are always the same globally, including their passwords. Consequently, you should not consider them as secure enough. Therefore, we advise you to delete them after you have added your own users.

Zoning methods

INRANGE provides both these types of zoning:

- ▶ Hard zoning
- ▶ Name server zoning (often called soft zoning)

Hard zoning makes it impossible to have I/O between different hard zones, whereas with soft zoning, one port could easily be assigned to any other port in the fabric.

However, both kinds of zoning are port related. It is the director's port that is referred to in both methods. At this time, a WWPN based zoning referring to the attached server or storage ports is not available. This will become available in the second half of 2001.

Both kinds of zoning are explained in 3.4.6, "Defining name server zones" on page 252.

The consequences are different for tape and for disk attachments. Most disk servers use LUN masking. This prevents the volumes from being used by another server other than the assigned one. So in an FC/9000 and ESS environment, unplugging the ESS cables from one zone and putting them into another zone, it will not be possible to access the ESS volume from servers that have not been associated with this storage.

However, this is different with tape. Tape drives normally do not support LUN masking. Therefore, the only access control is provided by the zoning methods of the director. Using the director-port based zoning concept, this tape drive could potentially be accessed by other hosts just by changing the cabling.

To limit the physical access to the director, consider some of these possibilities:

- ▶ Use the INRANGE cabinet with locked doors to prevent access to the cabling.
- ▶ Also limit the access to any patch panels used in your SAN infrastructure.
- ▶ Use access control methods for the datacenter facility.

Consequences of leaving the fabric unzoned

When the fabric with all directors and/or switches is in initial mode, then no name server zone is defined at all.

By default, all ports are in one hard zone called the default zone. Note that you always have to have all physical ports included in a hard zone. However, when using only name server table based zoning, you can just leave the hard zone setting as they are by default.

When leaving the fabric without any name server zone, the only effective zone is the all-inclusive hard zone. That means all ports are enabled to communicate with every other port.

Be aware of this any-to-any connectivity when leaving the fabric unzoned.

Once you have added one name server zone, then the nodes, if they are part of any zone, have limited access to the name server table.

Name server zone members can communicate with only those nodes that are part of the same name server zone.

3.3 Management of an INRANGE SAN

To gain access to all the management functions, we logged in to IN-VSN as a user with admin rights. In the topics that follow, we will explain how to:

- ▶ Set up operational parameters
- ▶ Connecting to an INRANGE fabric
- ▶ Assign names
- ▶ Define zones
- ▶ Cascade directors and merge fabrics

3.3.1 Setting up operational parameters

In this topic we describe the parameter setup which can be performed without having any directors known to the IN-VSN software:

- ▶ Manual database backup
- ▶ Automatic database backup
- ▶ Call Home and Page Home settings (CE's responsibility)

Database backup

INRANGE provides the ability to backup all director related data to an external disk. This is called database backup. You can use this backup to restore your configuration data. Database backup settings are performed using the IN-VSN server.

We will first perform a **manual backup**. Therefore, we use the IN-VSN server window to open the backup menu, as shown in Figure 3-23.



Figure 3-23 Manual backup of IN-VSN server: Step 1

Then we specify the name of our IN-VSN database file. We could leave the default name had we wanted to. However, we choose to have our own specific file named `itso_db2.mdb`, as shown in Figure 3-24.

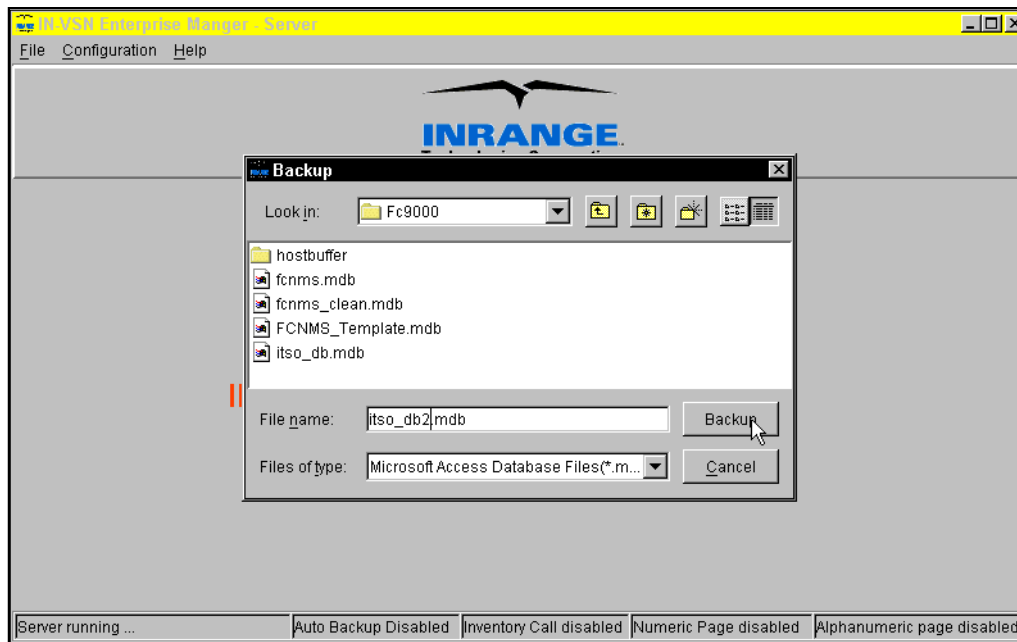


Figure 3-24 Manual backup of IN-VSN server: Choosing the filename

To achieve a higher level of disaster recoverability, we recommend that you place this database file on a mapped network drive, rather than use the management PC. By doing this, you can recover from the total destruction of the management PC. Another way of increasing recoverability is to perform periodic tape backups using backup software such as Tivoli Storage Manager.

The database file created by this manual backup could be used for IN-VSN database restore. However, in real life it may be difficult to perform manual backups periodically, since you need access to the IN-VSN server interface.

Therefore, it is possible to set up an **automatic IN-VSN database backup**. We recommend that you utilize this function, as it saves time, increases our recoverability, and does not rely on human memory to start the process.

The setup of autobackup is also done using the IN-VSN server interface. From the configuration menu, select **AutoBackUp** as shown in Figure 3-25.



Figure 3-25 AutoBackUp settings: Step 1

In the Auto Backup Settings window, activate the AutoBackupUp field as shown in Figure 3-26.




Figure 3-26 AutoBackUp settings: Enabling AutoBackUp

Again, we recommend placement of the backup file on a network device to achieve higher recoverability. We have decided to implement a backup interval of two hours instead of one hour as shown in Figure 3-27.




Figure 3-27 AutoBackUp settings: Specifying Filename and backup interval

The IN-VSN server main window then displays “Auto Backup Enabled” status.

Setting up page home/call home

In some environments it may be useful to have INRANGE or IBM automatically contacted via telephone or pager by the IN-VSN server. This is called Page Home or Inventory Update Call.

However these parameters should be set up by IBM or INRANGE Customer Engineers only.

An example of Inventory Update Call parameters is shown in Figure 3-28.

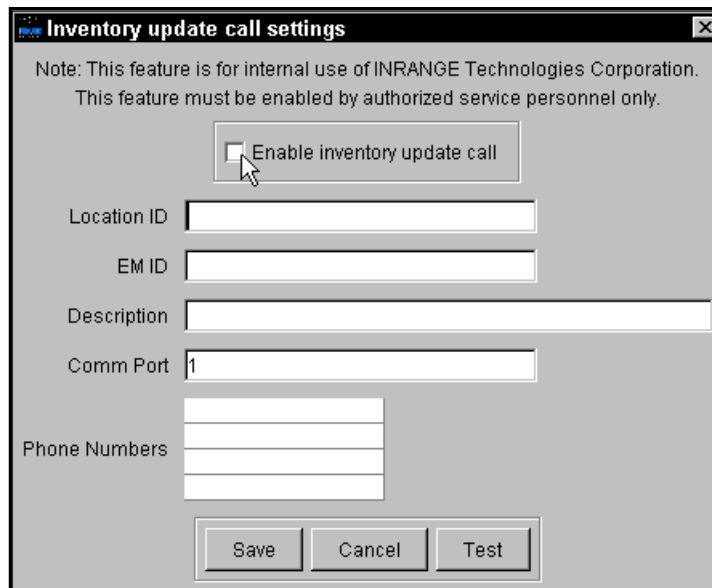


Figure 3-28 Inventory Update Call parameters

3.3.2 Connecting to an INRANGE fabric

A fabric is made up of one or multiple INRANGE directors and or switches which are linked via Inter Switch Links (ISLs). Directors that are not linked via ISL are considered as different fabrics.

Initially, the IN-VSN software is not aware of any fabric components at all, as shown in Figure 3-29

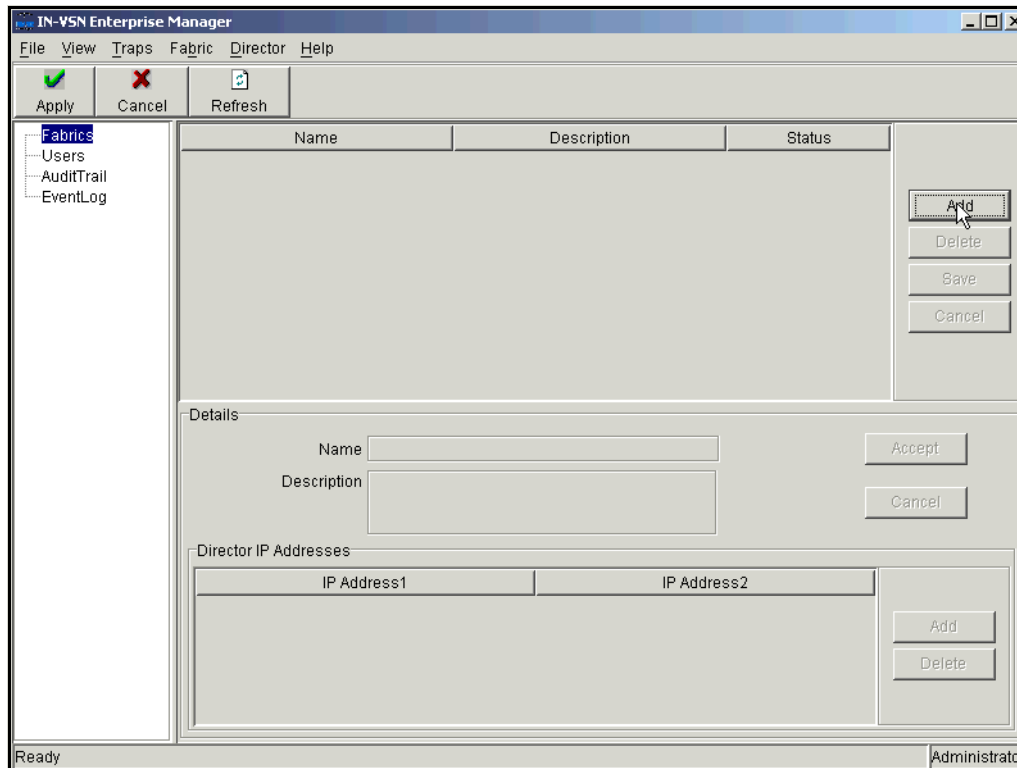


Figure 3-29 IN-VSN fabric tool without any fabric known to it

To make a specific fabric known to the IN-VSN software, click **Add**. Then enter all the required fabric information into the fabric details fields:

- ▶ We will call the fabric `osvl_fc64`.
- ▶ We do not give any description at this time.
- ▶ Our directors' IP addresses are 9.113.20.67 and 9.113.20.89.

For each director in the fabric, you must use a dedicated row in the Director IP Addresses field. In our case, we have one director, and therefore only one row. After entering all fabric information, we clicked **Accept** as shown in Figure 3-30.

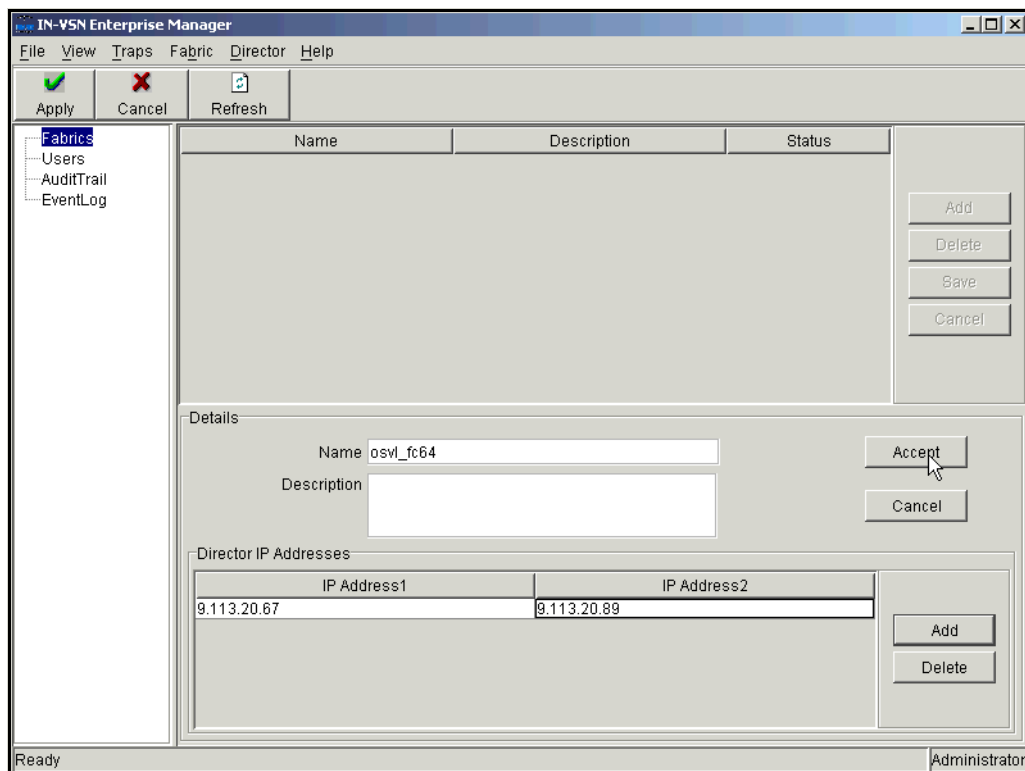


Figure 3-30 IN-VSN: Connecting to a new fabric

It is not necessary that the directors specified with their IP addresses be actually online or connected at the time of the fabric creation. The actual IP connectivity to these directors is obtained every time you click the specific **Fabric** tab in the IN-VSN client interface.

Furthermore, at the time of fabric creation, it is not checked if the specified directors are actually connected via ISL or not. Again, this will be obtained when clicking the specific fabric.

Once the director is actually online and IP connected it will appear by clicking the **fabric** icon as shown in Figure 3-31.

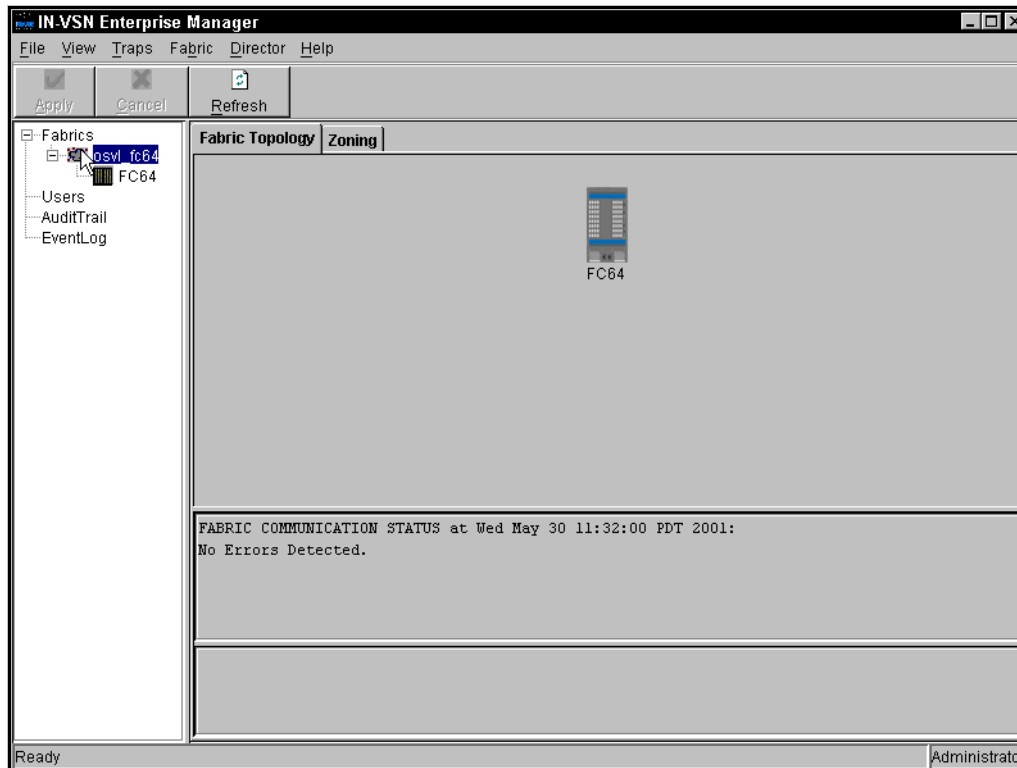


Figure 3-31 IN-VSN: Initial fabric view

The default name of this director is FC64.

If you had multiple directors in one fabric, they would appear as multiple symbols in the fabric view.

3.3.3 Assigning names and aliases

Our fabric consists of one director with 64 ports. We want to attach four Hewlett Packard (HP) servers and one Enterprise Storage Server (ESS) to this director.

Using the IN-VSN tool, we are able to assign names to physical director ports and to the director itself.

The assignment of names is not a requirement. However, we advise doing it because it will make our zoning and fabric management tasks much easier. This is especially true in larger environments with multiple directors.

To simplify further usage of this environment, we decided to use the names as shown in Table 3-2.

Table 3-2 Director and port names

Director component	Name to assign	Physical port of director
FC/9000 director	osvl_fc64	N/A
first port for D class server	osplsun2_td0	54
second port for D class server	osplsun2_td1	55
first port for L class server (HP1)	osplhp1_a	57
second port for L class server (HP1)	osplhp1_b	58
first port for L class server (GEODE)	geode_td0	60
first port for K class server	hewlett_td0	63
second port for K class server	hewlett_td1	62
ESS Bay1 Port4	ospl5b1p4	40
ESS Bay4 Port4	ospl5b4p4	48

These names belong to the physical ports of the director and not to the ports of attached devices.

Consequently, when changing the cabling of the director, you should also ensure that any change to the cable layout reflects the changes, and therefore the current environment.

Assigning the directors name

To assign the directors name we enter the director view by clicking the specific director in the Navigation explorer of the IN-VSN tool as shown in Figure 3-32.

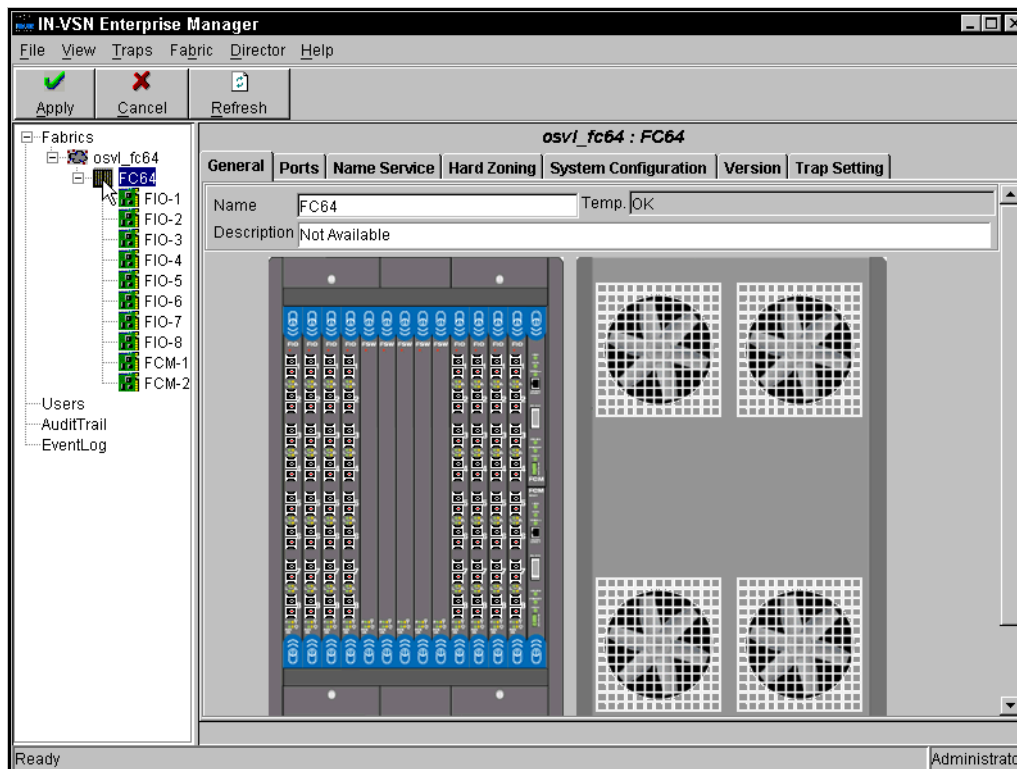


Figure 3-32 IN-VSN: Director view

We enter the new name `osvl_fc64` in the Name field in the director view. Click **Apply** and confirm this action to make the change effective as shown in Figure 3-33.




Figure 3-33 IN-VSN: Changing the directors name

Assigning port names

To assign names to individual FC ports of our director, we click the specific director in the IN-VSN navigation tree to have the director view displayed. Then we click the **Ports** tab to have a list of all 64 ports. This is shown in Figure 3-34.

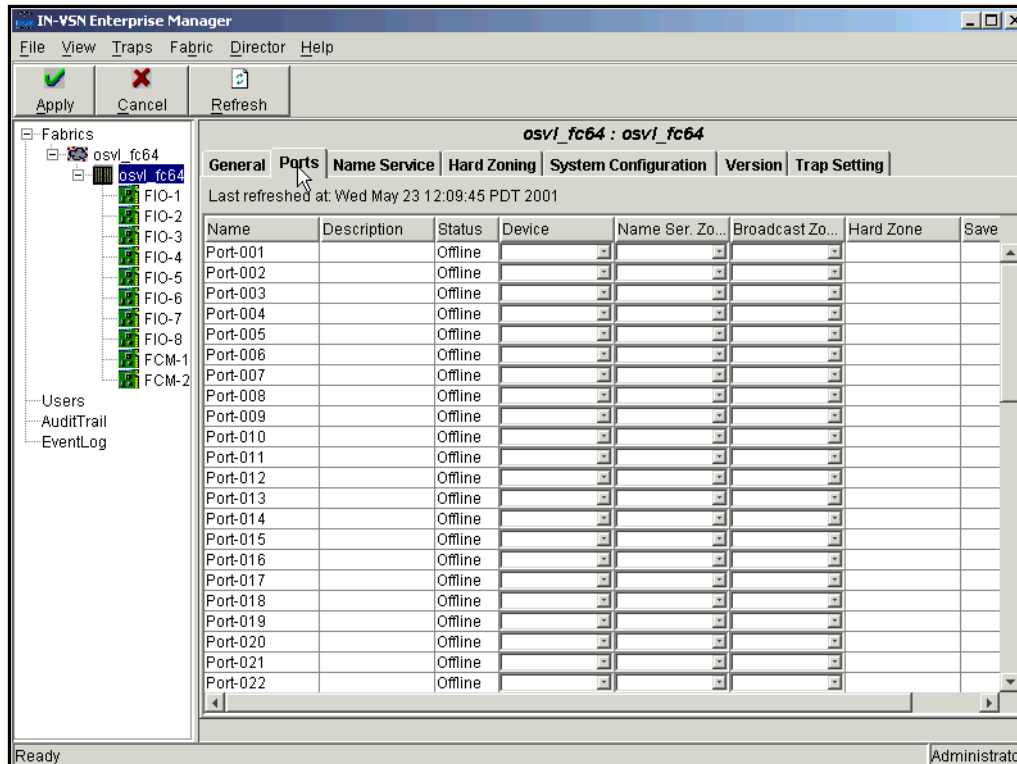


Figure 3-34 IN-VSN: Director view with all ports

This list displays all 64 physical ports of our director with their attributes, such as:

- ▶ Default names
- ▶ Status: whether you have signal light (Online) or not (Offline)
- ▶ Device: World Wide Node Name (WWNN) of the attached device or server
- ▶ Zones which the ports belong to

We can use this list to enter the new names in the Name column, as shown in Figure 3-35.

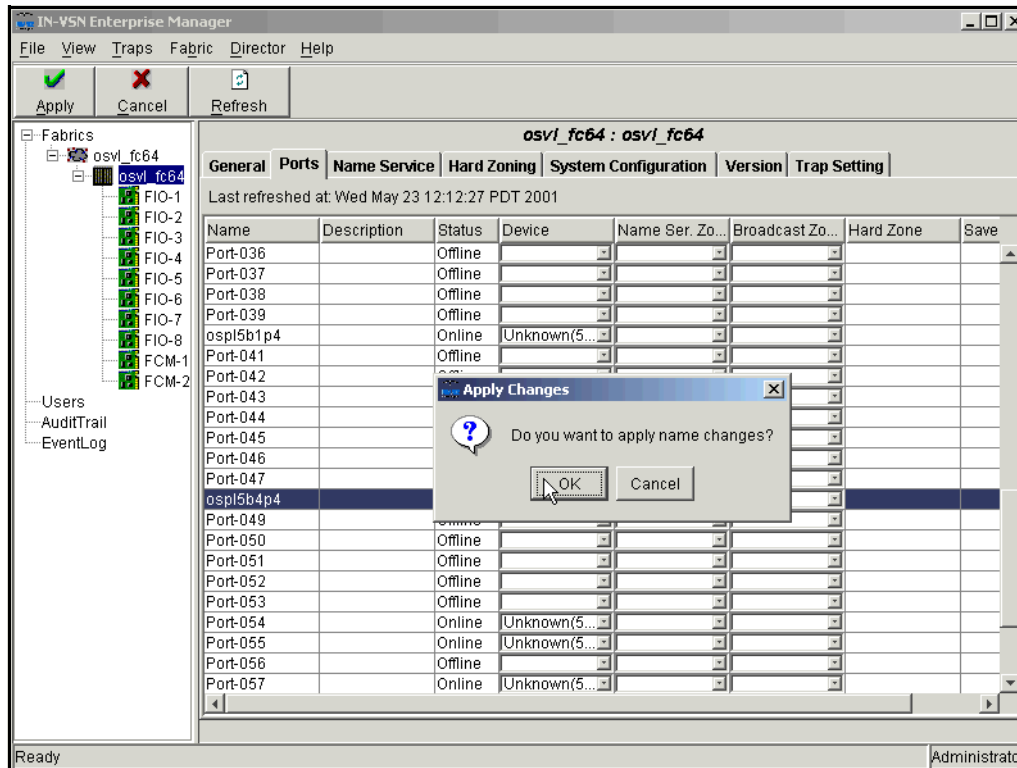


Figure 3-35 IN-VSN: Changing port names

We start with the two ports used for ESS attachment. Therefore, we change the name of port 40 to *ospl5b1p4* and port 48 to *ospl5b4p4*. Click **Apply** and confirm this action to make the new names effective.

Note: Take into account that all the port names you have assigned refer to the physical port of an INRANGE director. These names are technically independent of any attached external ports.

However, we advise you to use port names that indicate which external ports are actually attached or will be attached in the future. The port names can later on be used for zoning purposes. In case of zoning, using a meaningful port name will simplify the setup of zones. This is especially true for fabrics that span multiple director and switch chassis.

Then we repeat this procedure for all remaining server ports.

We have now physically attached all server ports and the two ESS ports. That is why we now have Device information (WWNN) displayed in the port list.

Figure 3-36 shows a port list with all of our desired name changes active and with the WWNN information displayed for the attached ports.

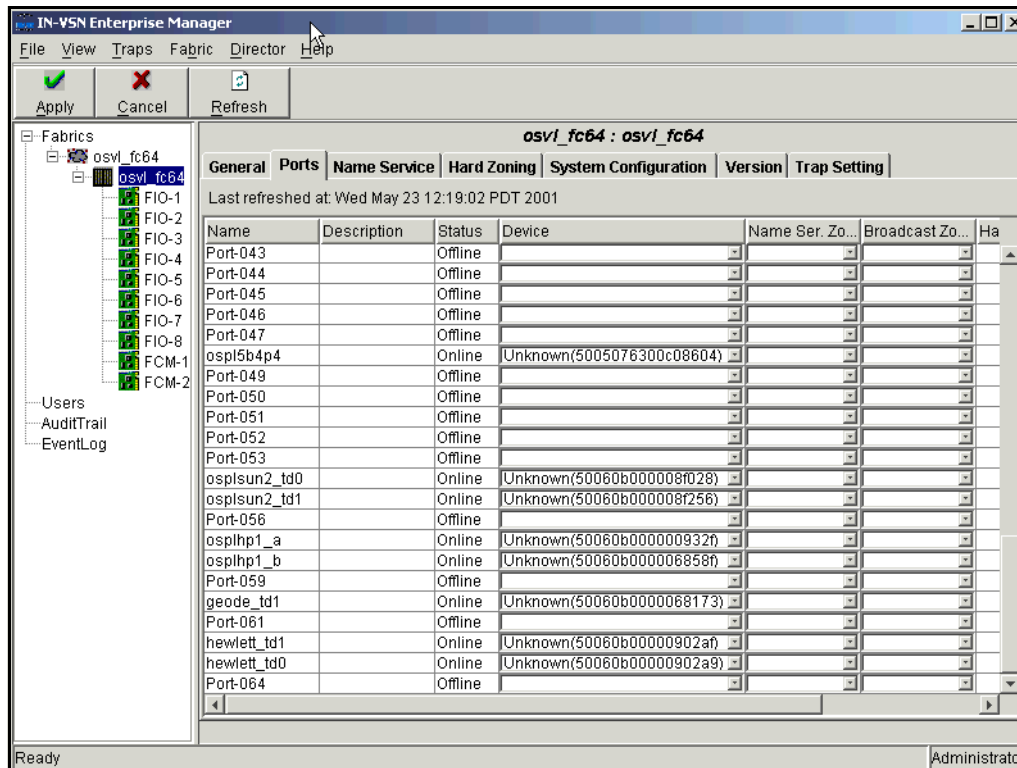


Figure 3-36 IN-VSN: Director's updated port view

Once the new names are effective they will be used throughout the whole IN-VSN fabric management, including zoning, FIO-blade monitoring and name services.

3.3.4 Setting the director clock

All directors will be delivered with preset time and date settings. However, in most cases, these clock settings do not match with the local ones.

These clock settings do not affect the fabric functionality at all. However it is important to set them because this makes reading and understanding the time-stamped logs much easier.

To set the director clock, just click the specific director in the navigation tree and then choose **Set Director Clock** from the **Director** menu, as illustrated in Figure 3-37.

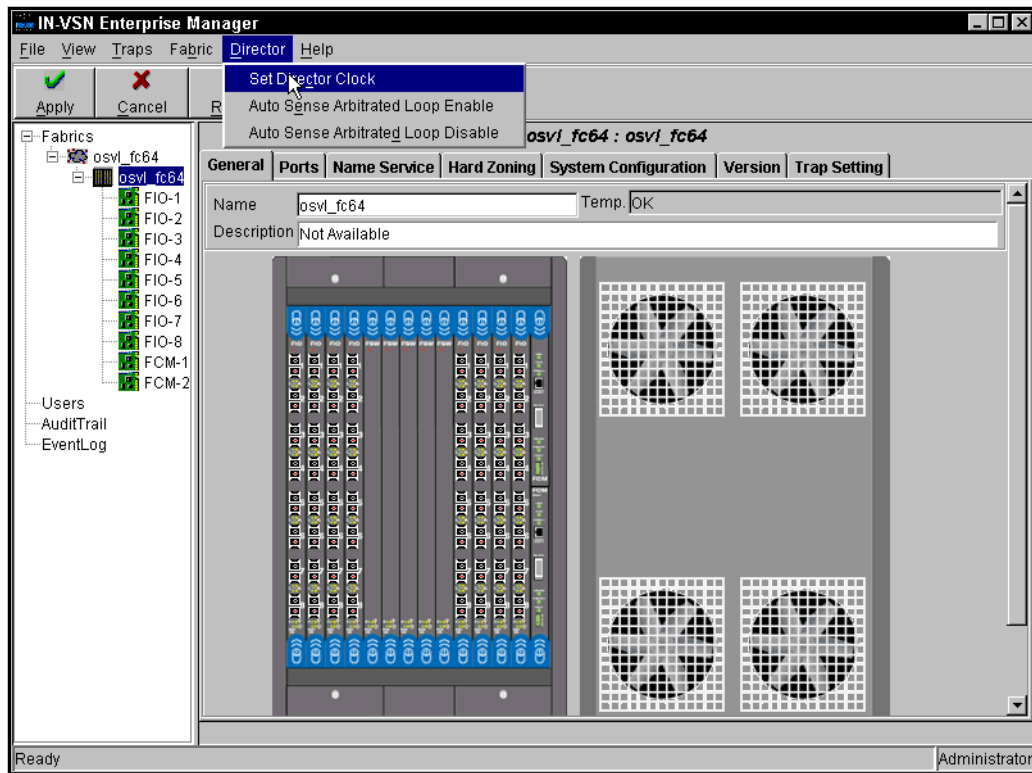


Figure 3-37 IN-VSN: Selecting the director clock menu

Enter your desired time settings and apply this by clicking **OK** as shown in Figure 3-38.

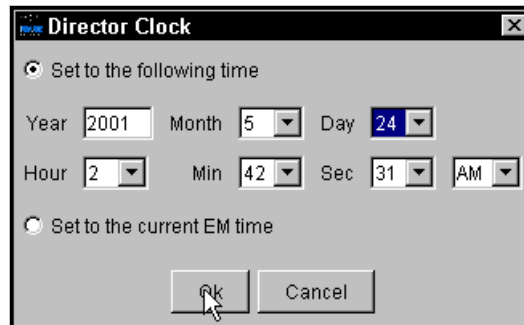


Figure 3-38 IN-VSN: Setting the director clock

3.3.5 Attaching loop ports

Today, the storage and server industry is moving rapidly towards switched fabrics. However, there are still a lot of systems that use Loop protocol. For instance, most tape devices use FC-AL, as well as lots of legacy FC host adapters.

Of course, some of them could be upgraded or replaced by newer devices to support native fabric login. However, for financial reasons it might be more attractive to protect these investments. This can be done using the FC-AL support features in an INRANGE fabric.

This topic covers:

- ▶ Enabling loop attachment for a specific director port
- ▶ Enabling loop attachment for the entire director
- ▶ Disabling loop devices
- ▶ Understanding differences between private loop and public loop
- ▶ Using Translative Loop ports to let private loop ports talk to the fabric
- ▶ Setting up private targets
- ▶ Setting up private initiators

Note: To read the following topics, it is useful to understand the differences between terms like director port, loop port, loop devices, initiators, targets:

- ▶ A director port is an actual physical port of the FC/9000 director.
- ▶ Loop ports are the external ports attached to a director port. Loop ports use loop protocols like private loop or public loop. They are sometimes referred as loop devices or loop nodes.
- ▶ Initiators are ports that actually control the IO flow and give commands. Typically initiators are FC host adapter ports.
- ▶ Targets are ports that get commands. Normally targets are storage devices such as tape drives or disk arrays.
- ▶ Loop devices running public loop are referred as NL-ports.

Enabling loop attachments

To make loop node attachment possible, we have to enable the director ports to autosense loop devices. We have two ways to make these settings:

- ▶ Enable Arbitrated Loop Auto Sense for the entire director
- ▶ Enable Arbitrated Loop Auto Sense for a specific director port

Enabling loop attachments for a specific director port

First we will allow loop attachment for the port called *hewlett_tdo*.

In the IN-VSN navigation tree, we select **FIO-8**, then the specific port named **hewlett_tdo**. In the port view we leave the **General** tab in the foreground. Here we activate the *Auto Sense Arbitrated Loop Enabled* field as illustrated in Figure 3-39.

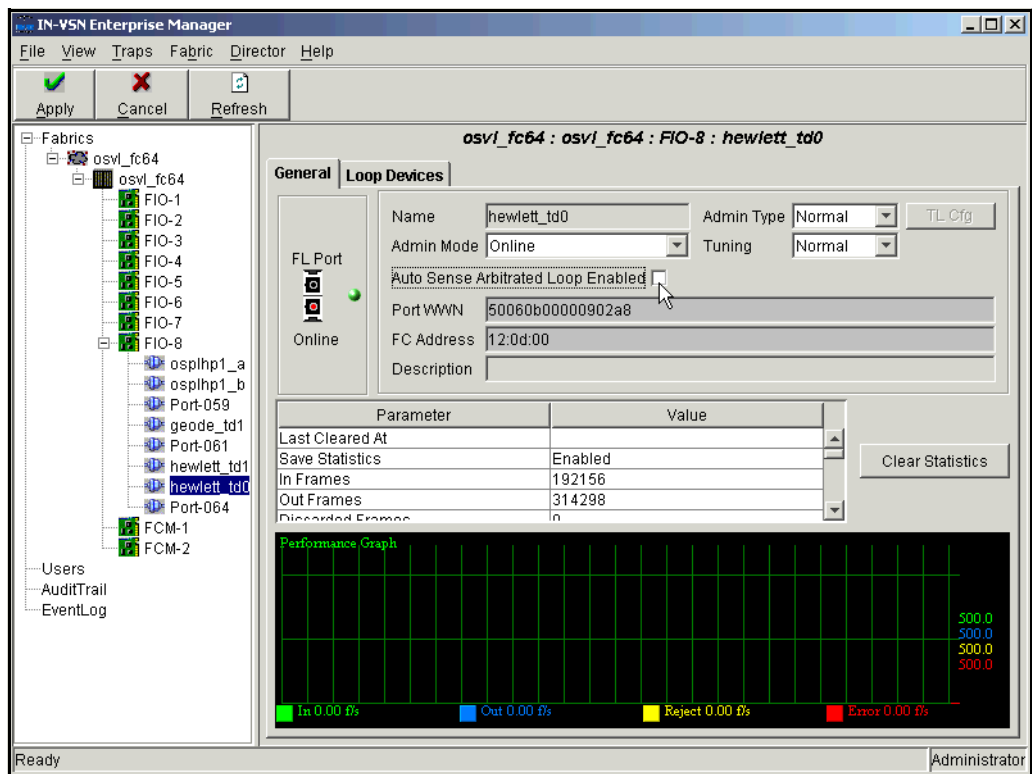


Figure 3-39 IN-VSN: Port View to enable port specific loop attachment

Once the Auto Sense field is activated in the port view, click **Apply** and confirm this action to make this change effective, as shown in Figure 3-40.

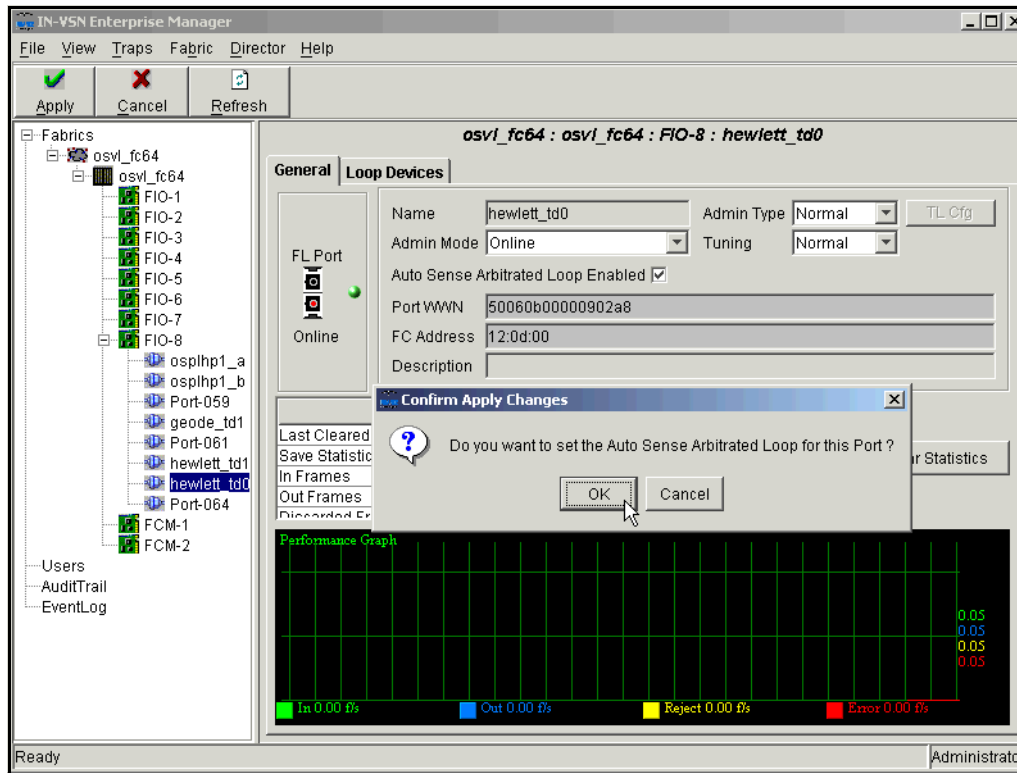


Figure 3-40 IN-VSN: Confirming the enabling of loop attachment

Now we are able to use loop attachments on this port. In our case, we have a Hewlett Packard Tachyon Lite Adapter set to **Public Loop**. The actual adapter settings are displayed in Figure 3-41.

```
# tdutil /dev/td0

Vendor ID is = 0x00103c
Device ID is = 0x001028
TL Chip Revision No is = 3.0
PCI Sub-system Vendor ID is = 0x00103c
PCI Sub-system ID is = 0x000006
Topology = PUBLIC_LOOP
Local N_Port_id is = 0x120d01
N_Port Node World Wide Name = 0x50060b00000902a9
N_Port Port World Wide Name = 0x50060b00000902a8
Driver state = ONLINE
Hardware Path is = 8/8/1/0
Number of Assisted IOs = 140841
Number of Active Login Sessions = 0
```

Figure 3-41 IN-VSN: FC adapter set to public loop

This adapter is set to **Public Loop** as shown in the Topology row.

By selecting the **Loop Devices** tab in the IN-VSN port view, we can ensure that the director actually recognizes this loop attachment as shown in Figure 3-42.

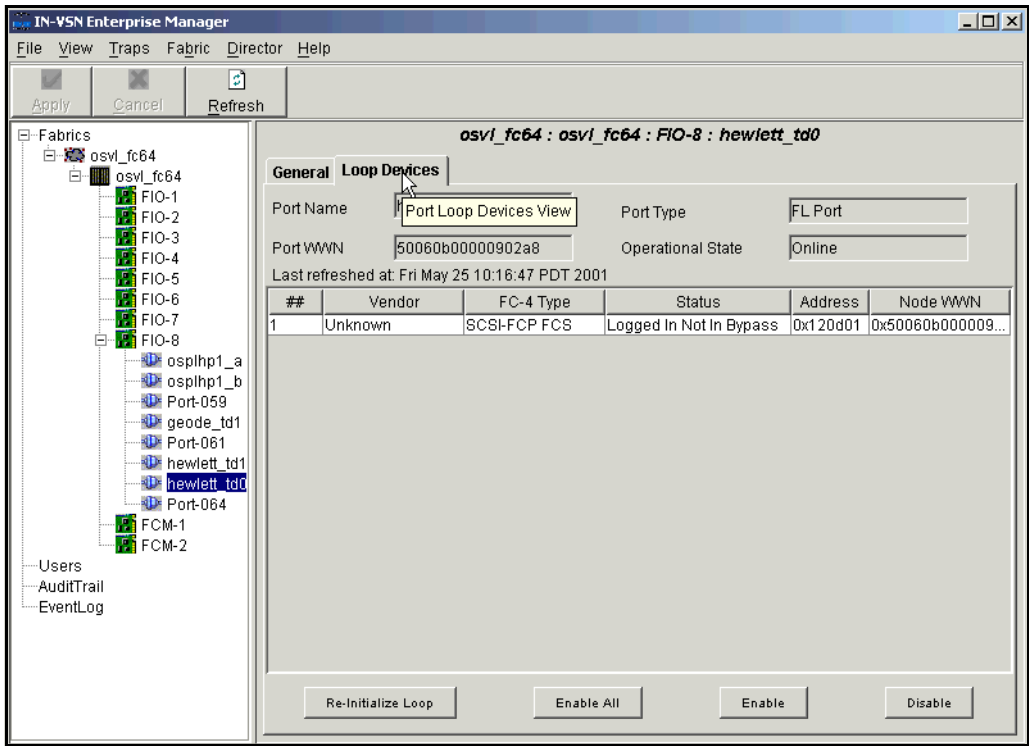


Figure 3-42 IN-VSN: Port Loop Devices View

Enabling loop attachment for the entire director

Instead of individually enabling loop attachment for lots of specific ports, you can enable this for the entire director as well. This means that all director ports will be set to *Auto Sense Arbitrated Loop*.

To enable all ports to attach loop devices, select the specific director in the navigation tree of IN-VSN and click **Auto Sense Arbitrated Loop Enable** in the **Director** menu, as shown in Figure 3-43.

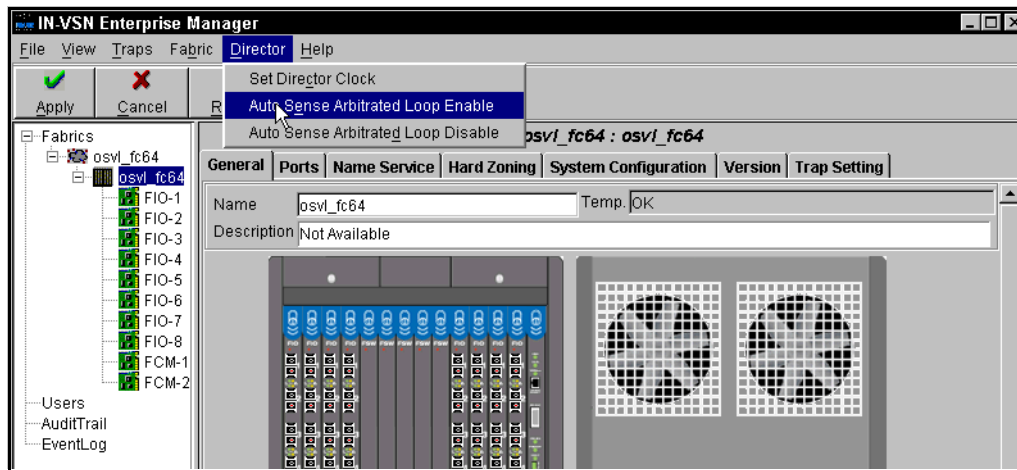


Figure 3-43 IN-VSN: Enabling Auto Sense Arbitrated Loop for the entire director

Be aware that this change could interrupt in-progress I/O on ports. Therefore, use this procedure only if you are sure about the consequences of the warning given in the confirmation window, as illustrated in Figure 3-44.




Figure 3-44 IN-VSN: Confirming loop enabling for the entire director

Loop ports in a name server table

Once you have a port enabled to sense Arbitrated Loop ports and there is actually a loop port attached, then you can verify that it has been recognized correctly by reading the name server table.

Public loop ports are displayed as NL_Ports (Node Loop), whereas private loop ports are displayed as TL_Ports (translative loop).

We show a representation of this in Figure 3-45.

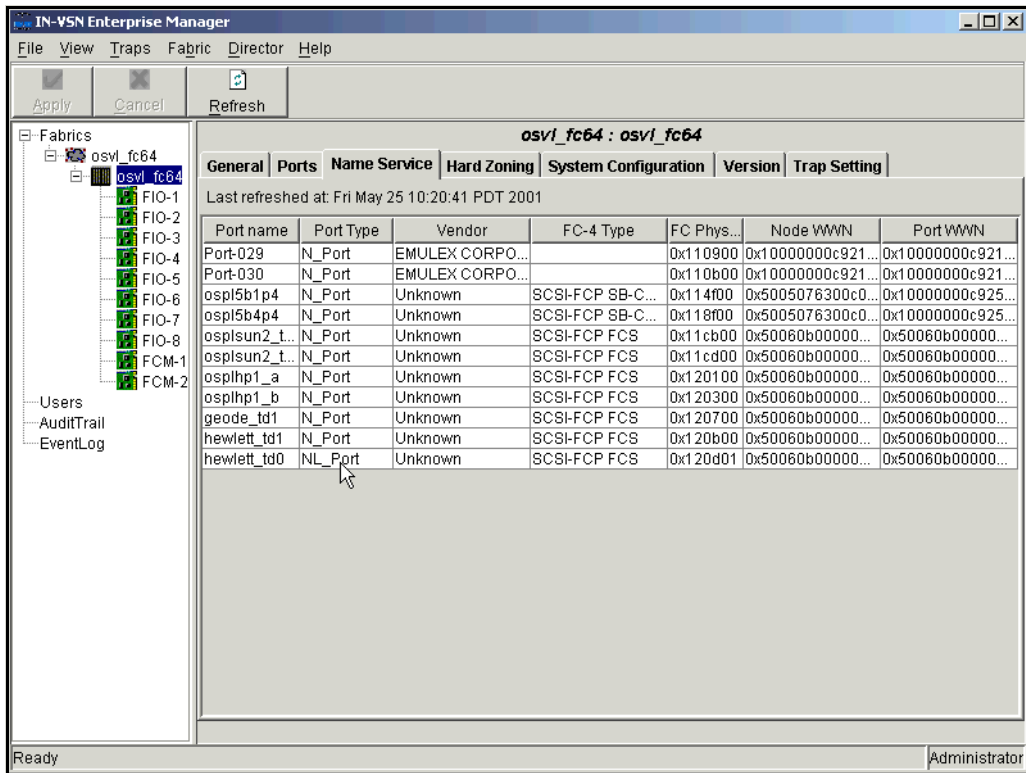


Figure 3-45 IN-VSN: loop port in name server table as NL_Port

Port *hewlett_td0* was set to enable AutoSense AL and we actually attached this port physically. After logging in, this port is displayed as an NL_Port, since this is a public loop port.

Bypassing loop devices

In cases where you have multiple loop devices attached to one director port, you can specify which devices should be actually used in the fabric.

In our example, we have four loop devices attached to one director port. This is shown in Figure 3-46.

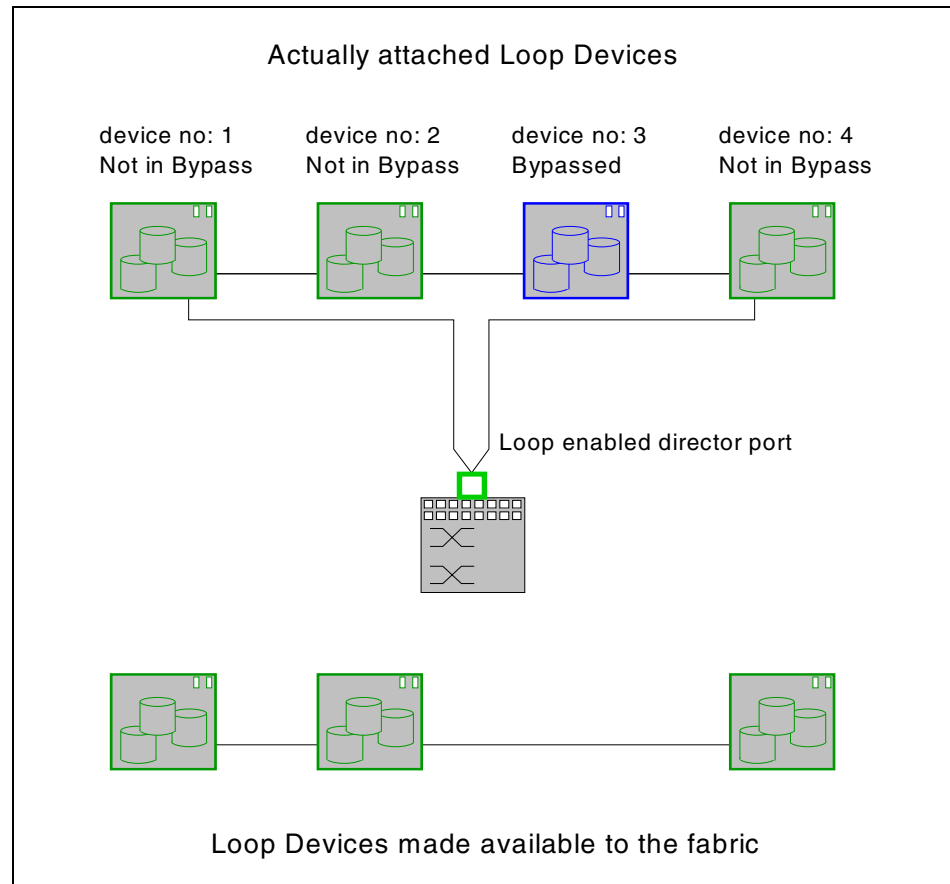


Figure 3-46 IN-VSN: Bypassing loop devices

We have decided that device 3 should not be seen in the fabric. To achieve this kind of filtering, device 3 must be disabled and thereby set to *Bypassed*. In this case, it means that a bypassed device is disabled from being used in the fabric.

This does not change the availability of the external Arbitrated Loop itself. For instance, even if a bypassed device fails it might affect the other external loop members as well. So, take into account that this kind of bypassing is just a way of filtering but not necessarily an improvement of the availability.

To disable a particular device select it in the Loop Devices menu of the specific director port. Click **Disable** to change its state to *Bypassed*.

To enable a particular device, use the same menu and click **Enable** as shown in Figure 3-47.

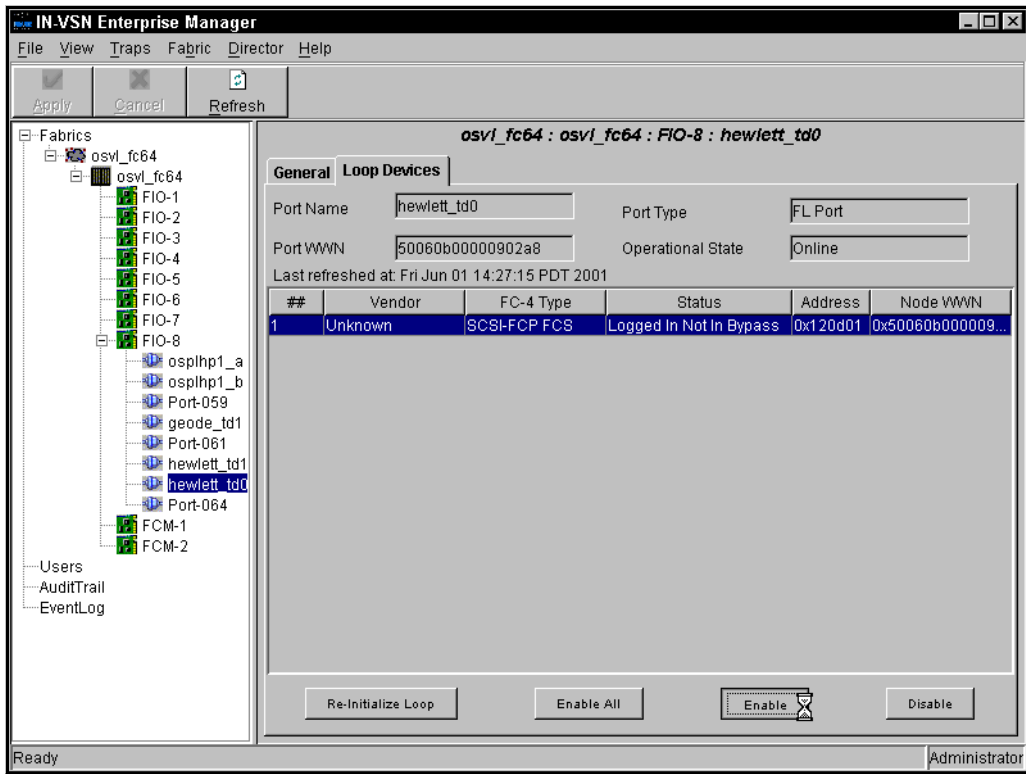


Figure 3-47 IN-VSN: Enabling and Disabling Loop devices

You should have at least one device per port left that is *Not in Bypass*. By default, all attached loop devices are set to *Not in Bypass*.

Note: Not all loop devices support Bypass. In such cases, even after clicking the **Disable** button, they will remain in *Not in Bypass* mode. That means such devices are always enabled.

Differences between private loop and public loop

There are two kinds of Arbitrated Loop protocols: *private loop* and *public loop*. For both kinds of loop ports, you have to enable the director ports to be able to use them as we have described previously.

However, there are some important differences between private loop and public loop, as discussed in the following topics.

Private loop

Private loop is the most basic form of Fibre Channel. Private loops are generally built using hubs, and communication is handled by passing messages around the loop from one device to the next in order to gain control of the loop and to communicate on it. Because each device is dependent upon the device next to it for communication, private loops tend to be unstable.

For example, adding, removing, or even rebooting a device causes other devices on the loop to pause. This pause forces every device to examine what changed in the environment and to adjust for that change. This *Loop Initialization Primitive Sequence* (LIP) can be extremely disruptive when data is being transferred at the same time that a change of state takes place on an arbitrated loop.

In a controlled, stable, and consistent environment, private loop can perform adequately for many applications. The main reason for its initial popularity is its relative low cost. However, when performance, stability and management are required, a higher performing Fibre Channel communication mode is necessary.

Importantly, a private loop is only known to itself. In some cases it is not possible, unless some form of translation takes place, for a private device to see or be seen by others.

Public loop

Public loop implies Fibre Channel arbitrated loop with hubs, switch or switches in place. Public loop adds a number of key advantages over private:

- ▶ **Name server** — As described in the Fibre Channel standard, the addition of a simple name server adds instant control and stability to a loop because each device registers with the name server on login. The name server then takes over the need to for devices to pause and pass messages to one another each a time an event happens on the loop. The switch internally controls the name server.
- ▶ **Management** — With a name server and switch management software, devices on a switch can be queried, examined, and brought online and offline. Errors can be logged and trapped with messages sent to an administrator who can then take appropriate action.
- ▶ **Zoning** — With zoning, a SAN can be made more secure by separating and segregating data and devices. Zoning can be done by individual switch ports, using name server entries, or by each device's unique world wide name (WWN).
- ▶ **Ability to address more than 126 Devices** — In a loop only configuration, it is possible to address a maximum of 126 devices, not including the initiator. Through connection to a fabric, that range can be increased. This is in comparison to a private loop's 8-bit address which limit a fabrics size.

Usage of public loop and private loop

Since these two loop modes inherit different protocol methods, there are differences in the actual usage of them with the INRANGE IN-VSN tool.

Usage of public loop

You can attach public loop initiators or targets without any additional settings, provided that the director port is set to enable AL.

However, attaching a public loop port does not automatically mean that this port can talk to any other ports in the fabric. This port will be recognized as an NL_Port and can be zoned in the same way as normal N_Ports can be zoned.

An example of the recognition of a public loop port as an NL_Port in the Name Service Table of a director is shown in Figure 3-48.

The screenshot shows the IN-VSN Enterprise Manager window with the 'Director' tab selected. The left sidebar shows a tree view of the fabric hierarchy, with 'osvi_fc64' expanded. The main pane displays the 'Name Service' tab for 'osvi_fc64 : osvi_fc64'. The table below shows the Name Service table with columns: Port name, Port Type, Vendor, FC-4 Type, FC Phys..., Node WWN, and Port WWN. The last row, 'hewlett_td0', is highlighted and has a mouse cursor pointing to it, indicating it is an NL_Port.

Port name	Port Type	Vendor	FC-4 Type	FC Phys...	Node WWN	Port WWN
Port-029	N_Port	EMULEX CORPO...		0x110900	0x10000000c921...	0x10000000c921...
Port-030	N_Port	EMULEX CORPO...		0x110b00	0x10000000c921...	0x10000000c921...
ospl5b1p4	N_Port	Unknown	SCSI-FCP SB-C...	0x114f00	0x5005076300c0...	0x10000000c925...
ospl5b4p4	N_Port	Unknown	SCSI-FCP SB-C...	0x118f00	0x5005076300c0...	0x10000000c925...
osplsun2_t...	N_Port	Unknown	SCSI-FCP FCS	0x11cb00	0x50060b000000...	0x50060b000000...
osplsun2_t...	N_Port	Unknown	SCSI-FCP FCS	0x11cd00	0x50060b000000...	0x50060b000000...
osplhp1_a	N_Port	Unknown	SCSI-FCP FCS	0x120100	0x50060b000000...	0x50060b000000...
osplhp1_b	N_Port	Unknown	SCSI-FCP FCS	0x120300	0x50060b000000...	0x50060b000000...
geode_td1	N_Port	Unknown	SCSI-FCP FCS	0x120700	0x50060b000000...	0x50060b000000...
hewlett_td1	N_Port	Unknown	SCSI-FCP FCS	0x120b00	0x50060b000000...	0x50060b000000...
hewlett_td0	NL_Port	Unknown	SCSI-FCP FCS	0x120d01	0x50060b000000...	0x50060b000000...

Figure 3-48 IN-VSN: Name server table with NL_Port

Usage of private loop ports

Normally, private loop ports are not able to talk with other ports outside of their own loop. Therefore, their addressing is different from normal ports. However, with the *Translative Loop* mode of specific director ports, it is possible to let them talk with other ports in the fabric.

To enable this you have to set the Admin Type of this particular director port to **TL** as shown in Figure 3-49.

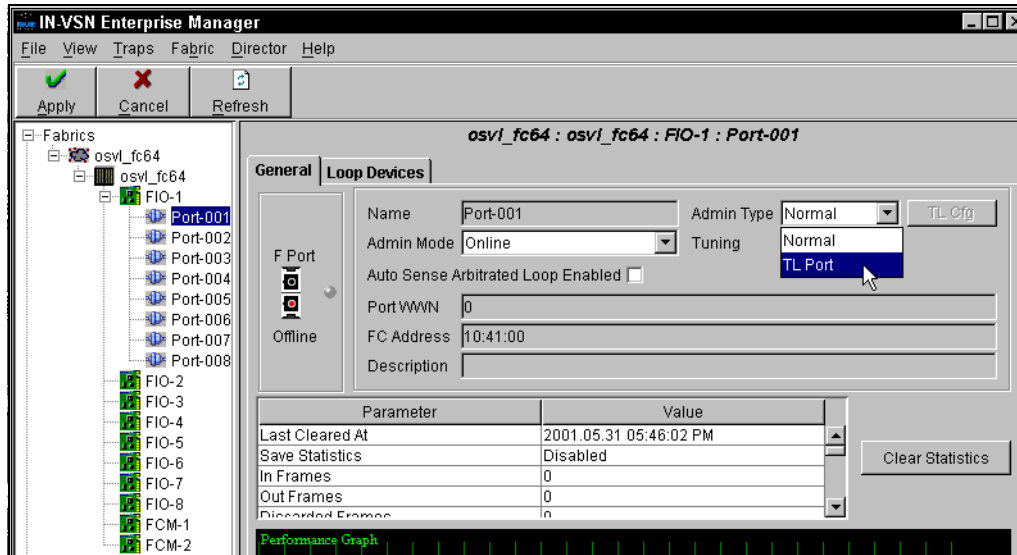


Figure 3-49 IN-VSN: Setting a port to Translative Loop mode

Each TL port has its own Translation Entries list. This list contains the set of off-loop devices which can communicate with this TL port. Off-loop devices can be:

- ▶ Any other N-port in the fabric (initiators or targets)
- ▶ Any other public loop port in the fabric (NL_Port)
- ▶ Any other private loop port in the fabric (TL_Port)

The creation and maintenance of this list is different between private targets and private initiators. Therefore, you have to specify whether a particular TL_Port has private targets or a private initiator attached to it. Click the **TL-CFG** button to open the TL Config window as shown in Figure 3-50.

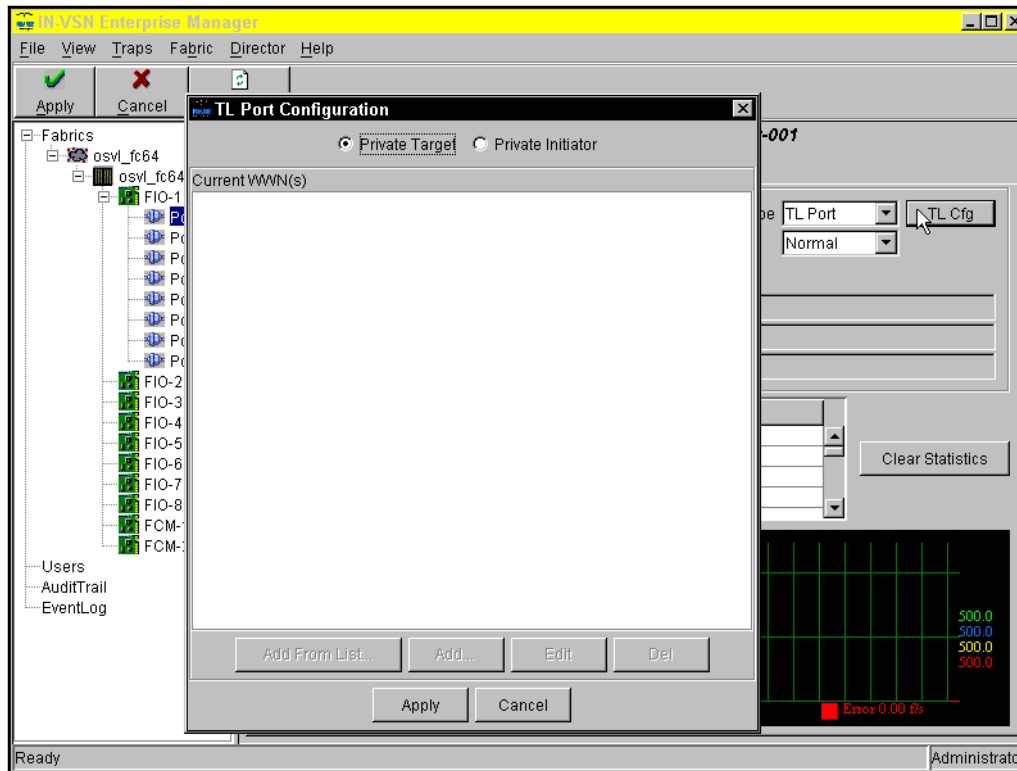


Figure 3-50 IN-VSN: Selecting target or initiator mode for TL ports

Select **Private Target** if you have targets attached to this TL_Port which should be accessed by other private initiators — for instance, private loop JBODs accessed by another private initiator.

Select **Private Initiator** if you have private loop initiators attached, such as elderly FC host adapters. Also, select this option if you want a private target device on this TL_Port to get accessed by a public initiator.

Note: The selection of Private Target or Private Initiator for a TL_Port does not necessarily mean that you have such types attached to it.

The private target option is only useful if a private initiator should access this particular target port.

If a public initiator should get access to this private target you should actually set this private target to Private Initiator. Then you should proceed to specify the WWN of the Public initiator that should have explicit access to this private target.

In such a scenario, the Private Initiator label for the radio button may be misleading.

The translation entries list for private target TL_Ports contains all initiators that try to get access to these targets. An Auto Learning feature enables this translation entries list to be updated automatically. That is why, for private targets, you cannot specify any WWNs in the TL-CFG window. However, you should limit the actual number of initiators trying to communicate with this TL_Port to a maximum of 31.

An Auto Learning feature for TL_Ports attached to private initiators is not available. All devices that this private initiator has to communicate with must be added to the list manually. Once you have selected **Private Initiator** in the TL-CFG window, some additional buttons become active. Click **Add from List** to add devices that are currently known to the fabric, as illustrated in Figure 3-51.

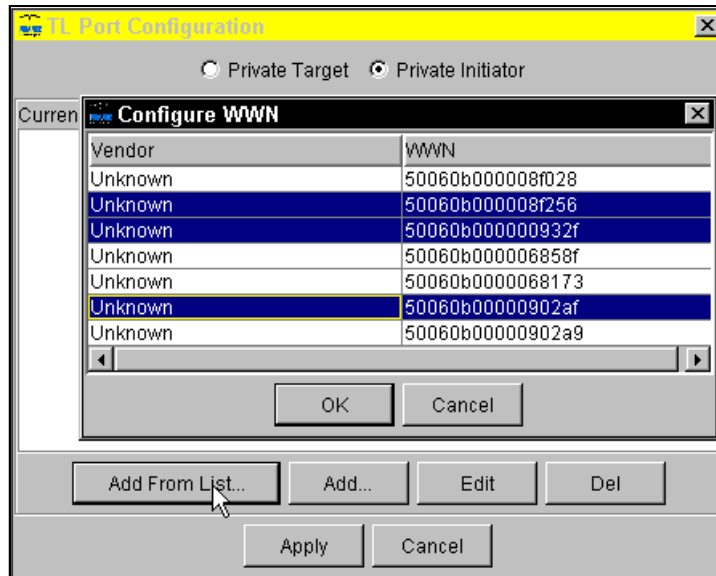


Figure 3-51 IN-VSN: adding possible targets

If you want to add device addresses that are currently not attached to the fabric, then you can add them by typing in their WWN. Click the **Add** button and enter the WWN as shown in Figure 3-52.

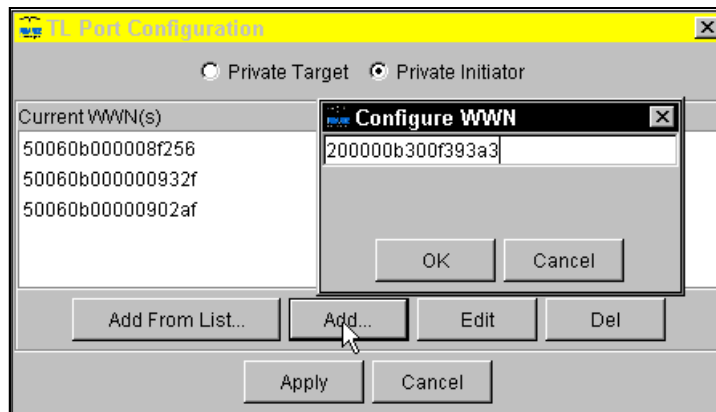


Figure 3-52 IN-VSN: Adding WWN targets

Each TL_Port has its own translation entries list, as shown in Figure 3-53.

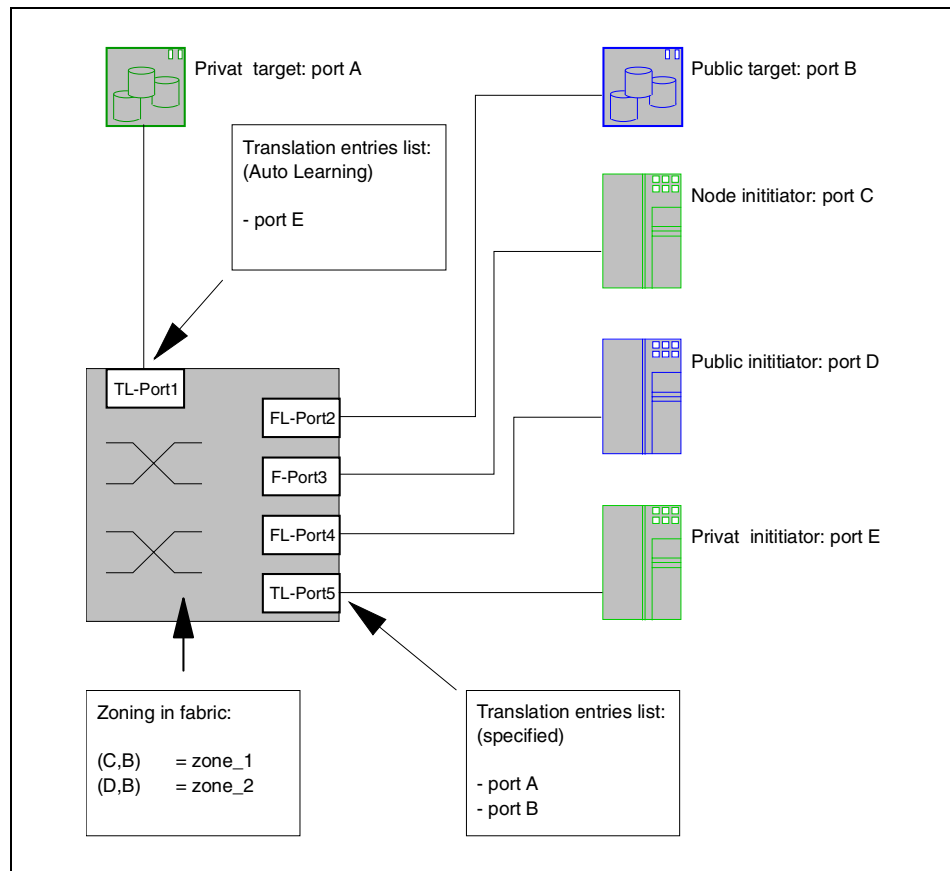


Figure 3-53 Translation Entries List with zoning

Only TL_Ports have translation entries lists, in our case port1 and port5.

The list of port5 (initiator) was manually created by adding the WWN of portA and portB. It means that port5 with its attached portE could have access to the ports A and B. Since TL-Cfg settings overrule name server zone settings, we do not need to zone these ports.

Furthermore, portD has access to portB. This is not implemented through a translation entries list, but by using zoning. Since portD is not a private loop port, there is no list for this port anyway.

The list for portA is automatically created using the auto learning feature. Once the initiator portE has actually logged in to portA, then portE will be added to the list.

Note: Even if the translation entries table for private target ports is updated automatically, you will not see its content. The auto learning feature is always used in the background. You cannot deactivate it or activate it.

TL_Ports and zoning

Once the translation entries list is created, these TL_Ports can actually talk to other ports in the fabric.

But the question remains: How does this TL_Port configuration interact with name server zoning, or with hard zoning?

First, TL_Port configurations cannot cross the boundaries of hard zoning.

Second, name server zoning does not have an impact on TL_Port communication at all:

- ▶ Private Target TL_Ports propagate all their translatable device addresses to the fabric. However, they can just be used by other Private Initiators. To enable communication between them the TL-CFG list of this private initiator must include the address of the private target.
- ▶ Private initiator devices have only access to devices specified in their TL-CFG list.
- ▶ Even if a TL port is part of a name server zone, it will not see the members of this name server zone unless they are put into the TL-CFG list.
- ▶ On the other hand, the TL-CFG list of a TL_Port is effective even if the name server layout would imply it otherwise.

Note that this feature of name server zones is limited to TL_Ports only. Normal public loop ports will be handled as normal director ports with all the normal consequences of name server zoning.

Impact of LIP in the fabric

Ports in loop networks use a process called *Loop Initialization Primitive sequence* (LIP) to establish their port addresses. All members of that loop are involved in a LIP.

Loop initialization occurs whenever there is a change in the layout of a loop, such as adding a new node, leaving of a node, or breaks in service in the loop.

The start of a LIP causes data transfers in progress to stop, thereby severely affecting the performance and availability of Arbitrated Loops.

However, since INRANGE is using a *Translation Entries List*, these LIPs will not be propagated to other fabric members. This is true even if multiple private loop ports are zoned together. Therefore, all LIP impact is limited only to the external physical loop (for example, an FC_AL hub).

This feature is a major benefit of using TL_Ports.

3.4 INRANGE zoning

One of the basic purposes of SAN fabric products is to enable or disable communication between the different ports (devices) attached to them.

In most cases it is helpful to limit the potential access of ports. Zoning provides an effective tool to limit and control the communication between fabric ports.

There are multiple possible reasons to limit access and include:

- ▶ We may want to avoid Windows servers seeing all disks in a fabric. Otherwise there would be a high risk of getting signatures written on all disks which would then mean these disks are unusable by other operating systems.
- ▶ For security reasons we want to limit the access to disk with confidential data to only selected servers.
- ▶ We would like to get control of the amount of paths a FC host adapter has to a specific disk. This is because not all environments are flexible in their usage of multipathing software.

INRANGE provides three kinds of zoning methods:

- ▶ Hard zoning
- ▶ Name server zoning
- ▶ Broadcast zoning

3.4.1 Understanding INRANGE hard zoning

Hard zoning follows physical boundaries within a single-stage switch chassis, and limits the communication of a port to only other ports in the same hard zone.

Hard zoning, in certain circumstances, is the only way to provide the required additional level of security, but careful consideration should be applied prior to activating any hard zones, as it may be possible to isolate devices.

- ▶ By no means is it possible to have communication over the boundaries of hard zones. This is also true if malfunctioning fabric initiators try to get around the name server tables.
- ▶ Hard zones take precedence over all other kinds of zoning (for example, over broadcast and name server zoning).
- ▶ If no hard zone is enabled at all then all ports are considered as being part of one large default hard zone.
- ▶ Name server zones and broadcast zones can be implemented within hard zones. If so, they further limit the connectivity between members of a hard zone.
- ▶ Hard zones can be created spanning multiple directors in one fabric.

Hard zoning rules

There are a number of rules that must be followed to implement hard zoning successfully:

- ▶ You can define a maximum of 16 hard zones in an INRANGE fabric, independently of how many chassis are used in the fabric.
- ▶ When a hard zone is created, it must be in a granularity of four ports.
- ▶ There is a fixed segmentation of a director into port groups which each have four ports. These port groups will be used when setting up hard zones.
- ▶ A single director port can only be part of just one hard zone. It cannot be a part of two hard zones at the same time.
- ▶ An *all-or-nothing* rule applies to hard zoning: Either *all* director ports are members of any hard zones, or *none* of them are members of hard zones.
- ▶ Any update to hard zone layouts will cause all members of affected zones to perform a fabric login. Hard zone changes should be restricted to initial setup and at maintenance slots.

Fixed placement of port groups

As mentioned before, the granularity to set up hard zones are groups of four ports. So one hard zone is built up of one or more multiples of these groups.

The location of these groups is fixed, and you cannot change this. A particular group consists of four ports: two ports on a FIO blade and two ports with the same location on a neighboring FIO blade as shown in Figure 3-54.

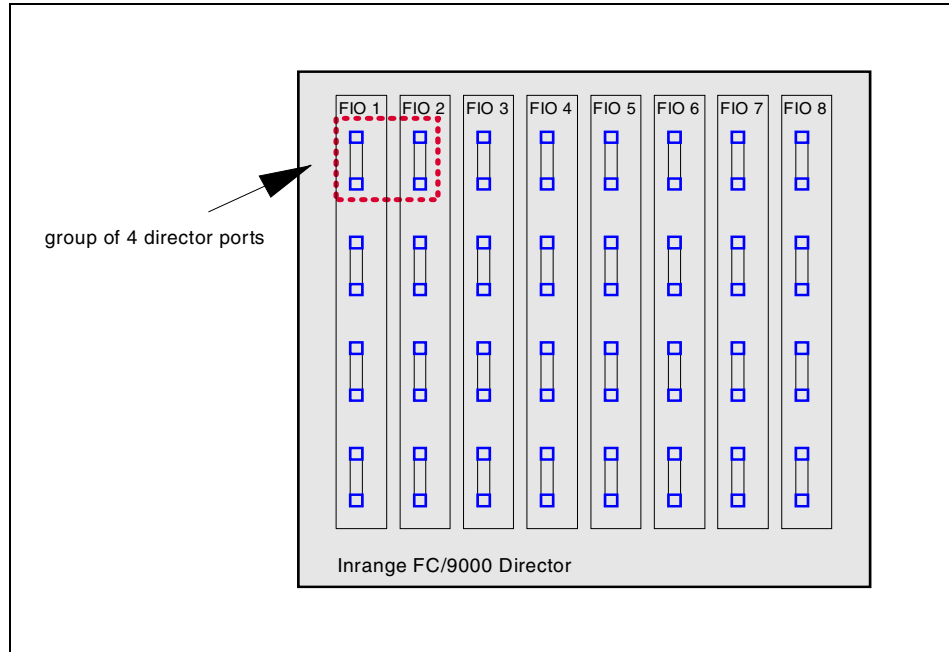


Figure 3-54 INRANGE hard zoning: fixed location of port groups

The two ports of one FIO blade being part of such a group are called adjacent ports.

The entire director is always automatically segmented into such groups. Consequently, with a fully equipped 64 port director, 16 port groups are automatically defined.

Building hard zones using port groups

Now we know that a hard zone must consist of port groups. In our example, we have created two hard zones, as illustrated in Figure 3-55.

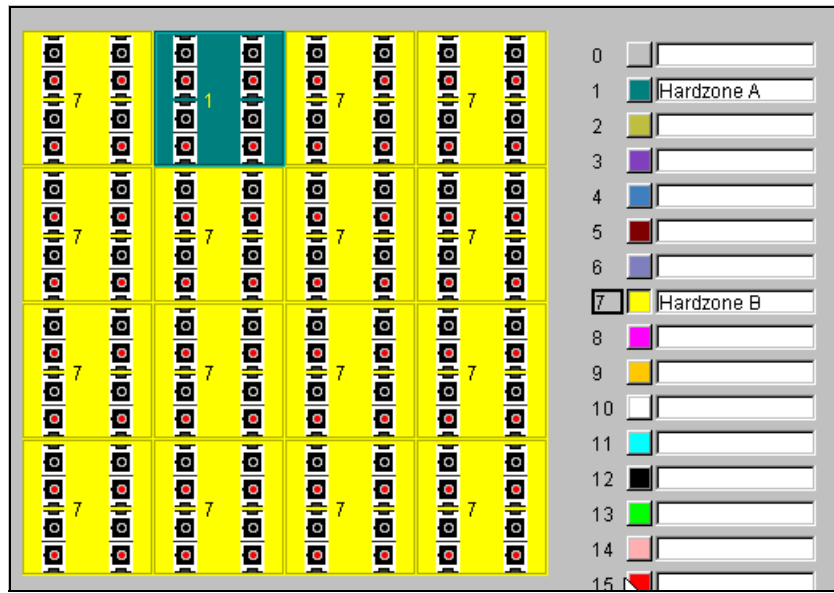


Figure 3-55 INRANGE hard zoning: basic example with 2 zones

This picture was taken by using the IN-VSN tool.

Hard Zone A consists of only one port group. Hard Zone B consists of the remaining 15 port groups. Remember the “all-or-nothing” rule: When implementing hard zoning, all ports must be zoned.

In the example above, the four ports of Zone A are allowed to talk to each other. All 60 ports of Zone B are allowed to talk to each other. However, any communication between Zone A and Zone B is blocked.

Remember that hard zoning strictly excludes ports other than those in the same hard zone from communicating together.

When creating hard zones, all port groups in a particular hard zone must be adjoining. In Figure 3-56 we show an example of an incorrect implementation.

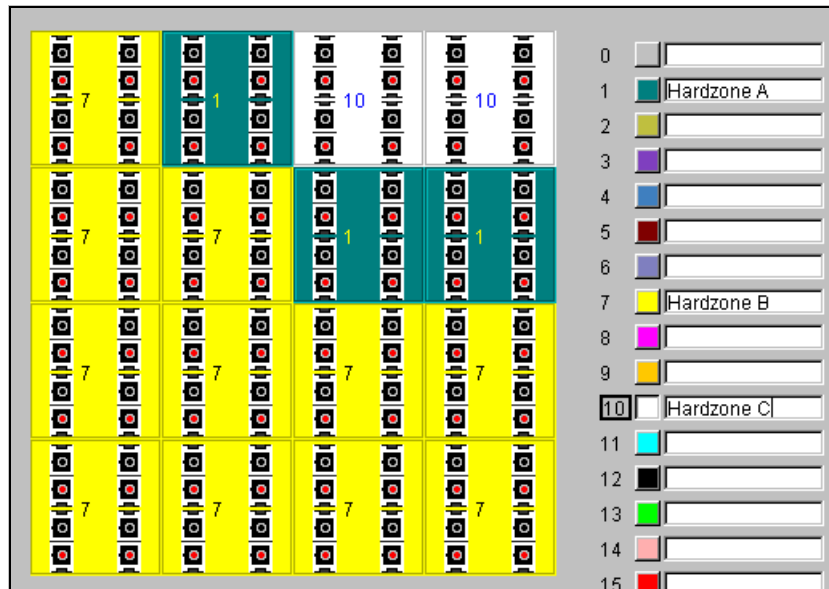


Figure 3-56 Violating the adjoining rule

The upper port group of Hard Zone A has no adjoining contact to the remaining two port groups of Hard Zone A. This is a violation of the hard zoning rules, and therefore it is not possible to create a hard zone layout such as this.

3.4.2 Understanding INRANGE broadcast zoning

Broadcast zones focus the distribution of broadcast messages to only those targets that need to receive them. This results in reduced fabric traffic and eliminates unnecessary message processing. Broadcast zones can overlap and are also assigned by director ports.

The granularity in which broadcast zones are built is one director port.

However, broadcast zoning is used for IP networking. Storage area networking environments do not exploit this kind of zoning.

3.4.3 Understanding INRANGE name server zoning

In contrast to hard zoning, there is no actual physical segmentation of fabrics with name server zoning.

Instead, a *name server table* is used to implement that type of zoning.

Name server tables store information about nodes that have logged into the fabric. Name server zones restrict the access of affected ports to this information. All unzoned nodes have full access to the name server table.

This means that zoned ports logging into a fabric will only get name server table information of other ports which are part of the same name server zone. Unzoned ports are considered as being part of an *orphan zone* which is not visible to users, but allows communication for all unzoned nodes.

Because this kind of zoning relies on the proper usage of FCP protocol commands, it is also widely called *soft zoning*.

Name server zoning gives better flexibility than hard zoning:

- ▶ A particular director port can be part of multiple name server zones.
- ▶ The granularity of ports to build a name server zone is only one port.
- ▶ Name server zone members need not be physically adjoining each other.
- ▶ The change of name server zones does not enforce a port re-login of affected ports.

However, since name server zoning relies on the correct usage of FCP commands, there is still a risk that malfunctioning nodes would affect other ports even if these other ports are not members of the same zone.

Due to its strength and flexibility, name server zoning is widely used in open systems environment.

Name server zoning rules

There are rules that must be adhered to:

- ▶ INRANGE name server zoning refers to physical director ports. A type of name server zoning pointing to WWNs of attached nodes will be available in the future.
- ▶ Each name server zone needs to get a unique number and a name.
- ▶ As many as 256 zones are possible in a fabric.
- ▶ Name server zones cannot cross the boundaries of defined hard zones.
- ▶ Director ports which have private nodes attached to it need to be set to TL mode. name server zoning is not effective for TL_Ports.

3.4.4 Hard zones and name server zones together

Hard zoning can strictly separate port groups, and this can be seen as an effective security feature.

Name server zoning allows us to further define the communication control on a per port granularity.

When name server zones and hard zones are used in conjunction, we need to consider the following principles:

- ▶ If you have no hard zone(s) created at all, then all director ports are considered as being part of one big default hard zone.
- ▶ However, if at least one hard zone is implemented then this default hard zone is not effective anymore. So, be aware of the fact, that if you add just one hard zone all remaining ports must be added to a hard zone too. For instance you could create a second hard zone containing all the remaining ports.
- ▶ Multiple name server zones can exist within one hard zone.
- ▶ Name server zones cannot cross the boundaries of hard zones
- ▶ Using INRANGE we do not have to worry about zone sets or active and inactive zones. Creation of zone sets or creation of passive zones is not possible in INRANGE environment. Either a zone exists and is thereby active, or it does not exist.
- ▶ All nodes that are not part of any name server zone have unlimited access to the name server table. This is also true if no name server zone is implemented at all. Consequently, by default, all attached nodes have access to the name server table. Only those nodes that are part of name server zones will have limited access to this information
- ▶ The only exception is a TL_Port Config list which explicitly allows access to the specified TL_Port even if name server zones exist that normally would imply otherwise.

So, we can see that we have several different ways to control actual node access with INRANGE:

- ▶ Hard zoning
- ▶ Name server zoning
- ▶ Translation Entries lists for TL_Ports

The following rules are enforced with INRANGE directors:

- ▶ Both TL-CFG lists as well as name server zones cannot span hard zone boundaries.
- ▶ You can use name server zoning to further limit access between nodes
- ▶ TL-CFG lists are for TL_Ports only. They overrule the name server principles. So, TL_Port attached nodes can only communicate with ports specified in the TL-CFG lists.

This hierarchy is shown in Figure 3-57.

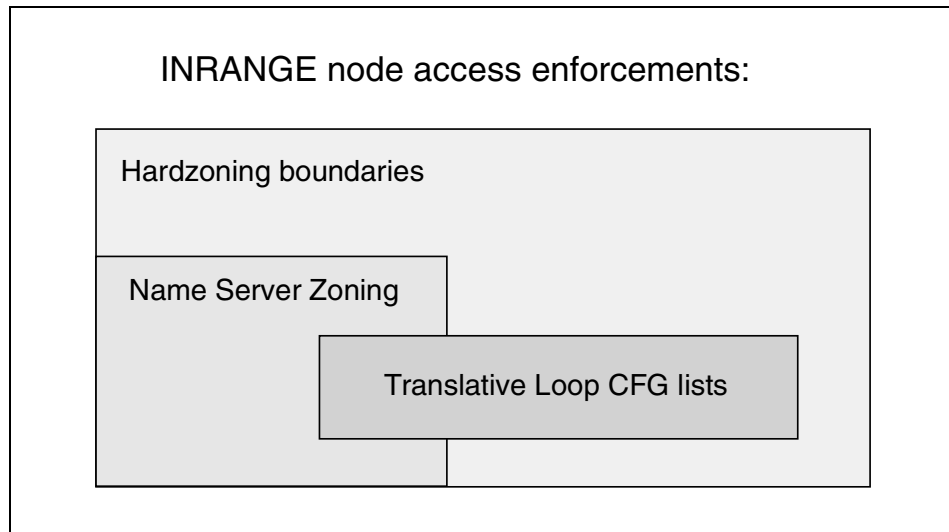


Figure 3-57 INRANGE access enforcements

The scenario illustrated in Figure 3-58 describes a possible layout.

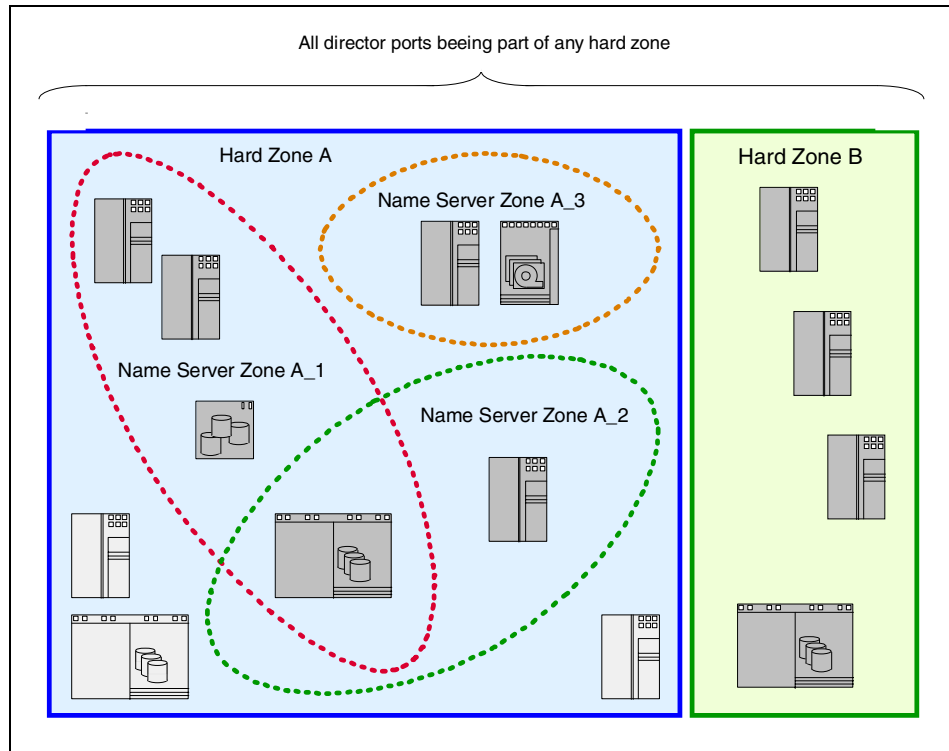


Figure 3-58 INRANGE fabric scenario with hard zoning and name server zoning

In this layout:

- ▶ All three defined name server zones are within Hard Zone A. They cannot span over to Hard Zone B.
- ▶ Name server zone A_1 and A_2 are overlapping. Both have an ESS node as a member.
- ▶ All members of name server zones can only access other members within the same Name server zone.
- ▶ All remaining nodes in Hard Zone A which are not part of any name server zone are part of the *orphan zone*. This allows communication between these three nodes.
- ▶ There is no name server zone at all in Hard Zone B. All members can have access to any other member in Hard Zone B since they are put into the orphan zone. However, this orphan zone is limited to Hard Zone B.

3.4.5 Defining hard zoning

In our example shown in Figure 3-59, we want to separate eight director ports for usage by the finance department only. In no way should it be possible for anyone else to gain access to their ports. That is why these eight director ports will be put into one dedicated hard zone (the Finance hard zone).

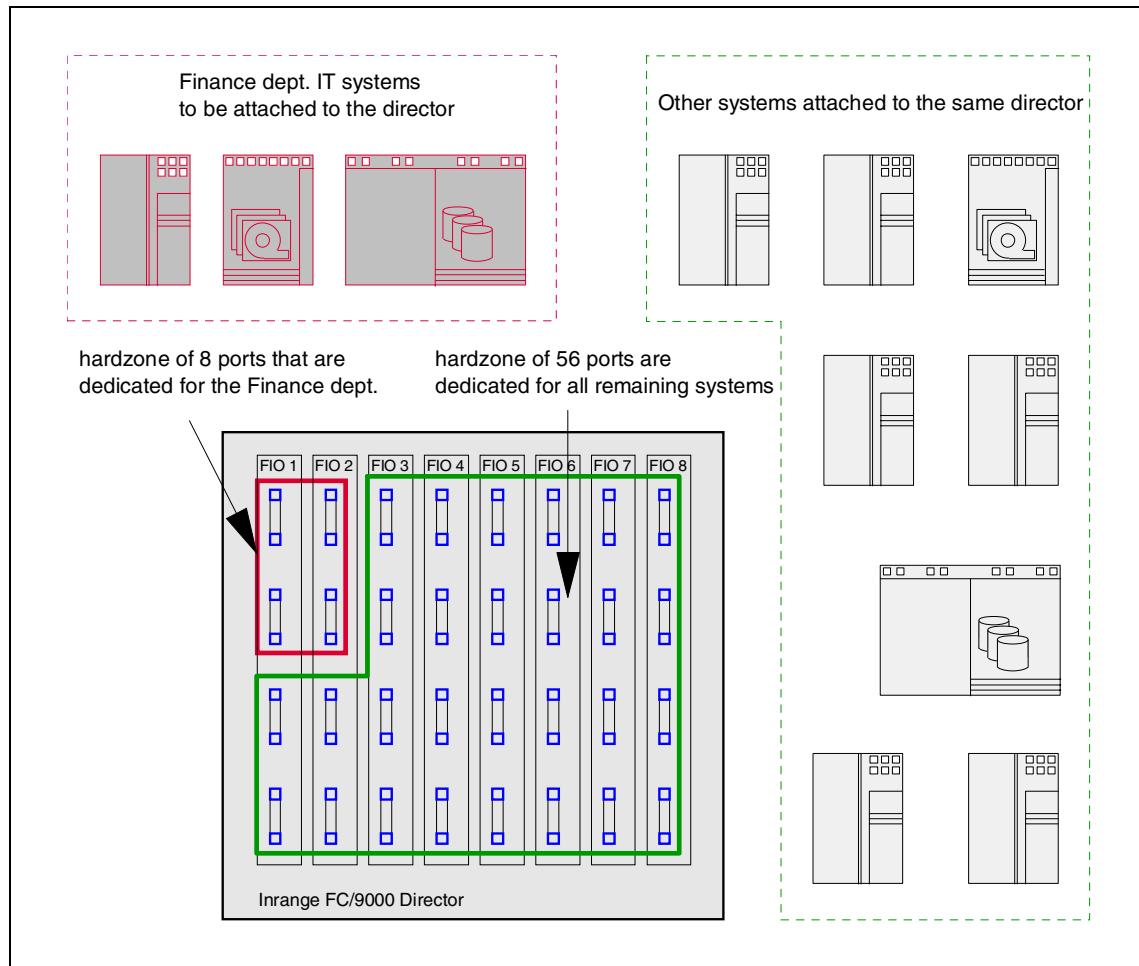


Figure 3-59 INRANGE hard zoning: Layout scenario

Considering the hard zoning rules that we have already discussed, we will place all remaining ports into a second zone. This second zone can be used by all other systems. Name server zoning is possible in both hard zones.

To start our hard zoning setup, we first select the particular director in the navigation tree and then click the **Hard Zoning** tab as shown in Figure 3-60.

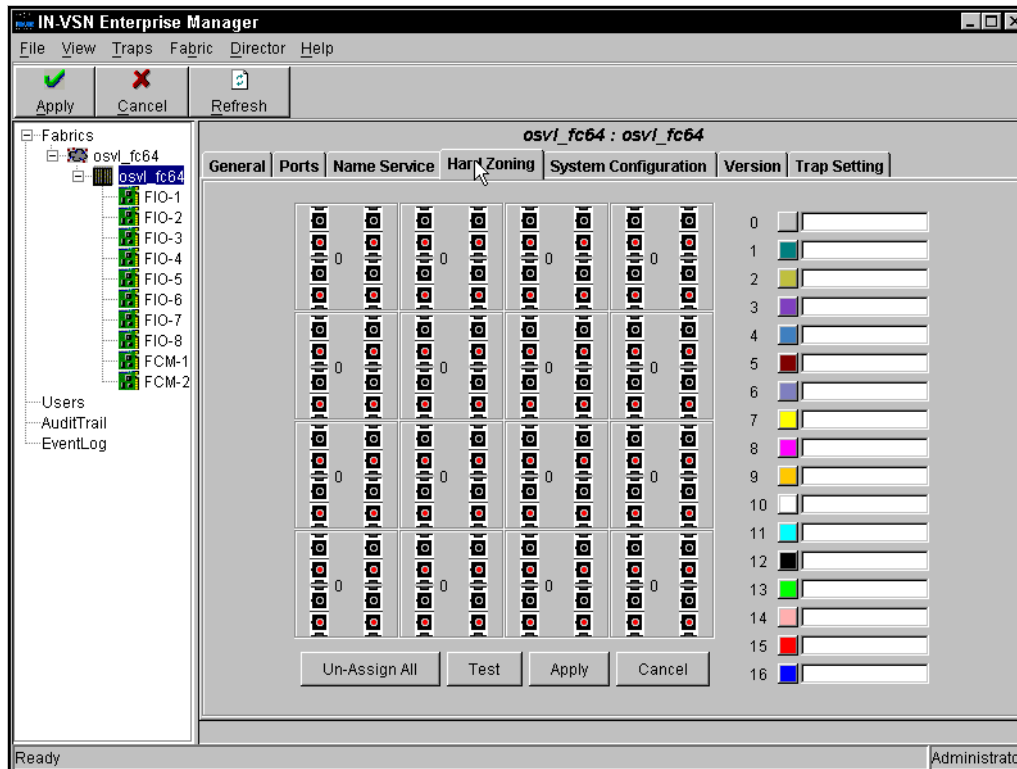


Figure 3-60 IN-VSN: Selecting hard zoning in the director view

This window is used for all hard zoning configurations related to the selected director.

This include actions such as:

- ▶ Renaming a zone
- ▶ Adding a zone
- ▶ Deleting a zone
- ▶ Change the zone layout

To create a specific hard zone, select one of the colored rows on the rightmost side of this window and type in a name for this zone. This is illustrated in Figure 3-61.

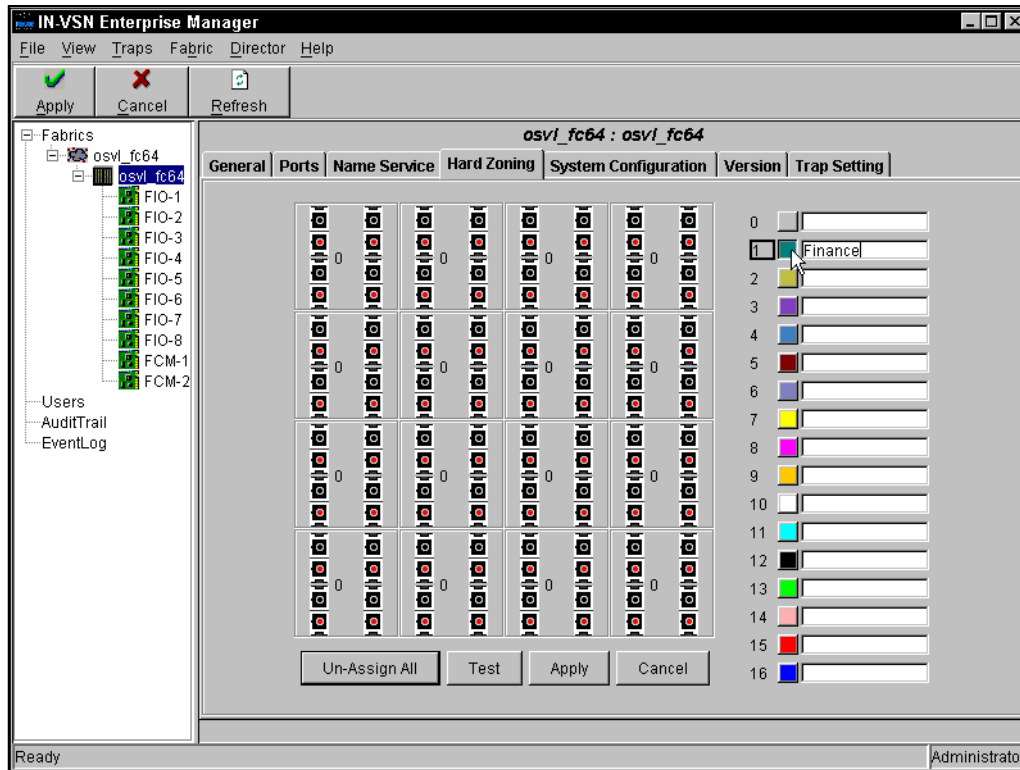


Figure 3-61 IN-VSN: Specifying a name for a hard zone

We have chosen to name the first zone *Finance*. However, later on, we could change the name again.

The remaining ports will be placed in a second zone, which we will call *Other*.

Note: The used color or number for a dedicated zone does not have any effect on the actual behavior or performance. These numbers and colors are solely used to simplify the organization and usage of hard zoning.

Hard zone colors, names or numbers are not visible in any way by attached nodes.

To include eight ports as members in our first zone, *Finance*, we click the port groups we want to use while the *Finance* zone is still marked in the rightmost table.

We have selected the two upper leftmost port groups to be part of hard zone **Finance** as illustrated in Figure 3-62.

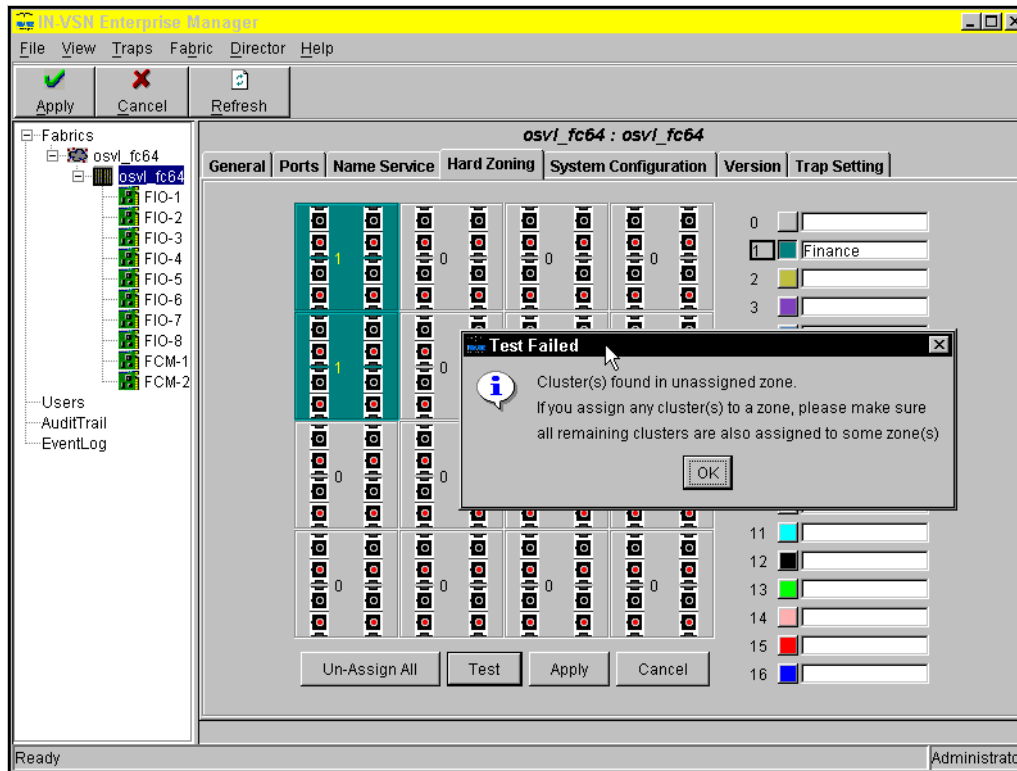


Figure 3-62 IN-VSN: two ports assigned to a hard zone

To test the integrity of our zone layout, we click the **Test** button. As you can see, we got a **Test Failed** feedback.

This is because we violated a hard zoning rule: Either all ports of a fabric belong to hard zones or none of them are hard zoned.

Consequently, we will place all remaining director ports in a second zone called *Other*, as we have described before.

Again we test this layout by clicking the **Test** button. This time everything looks fine, as illustrated in Figure 3-63.

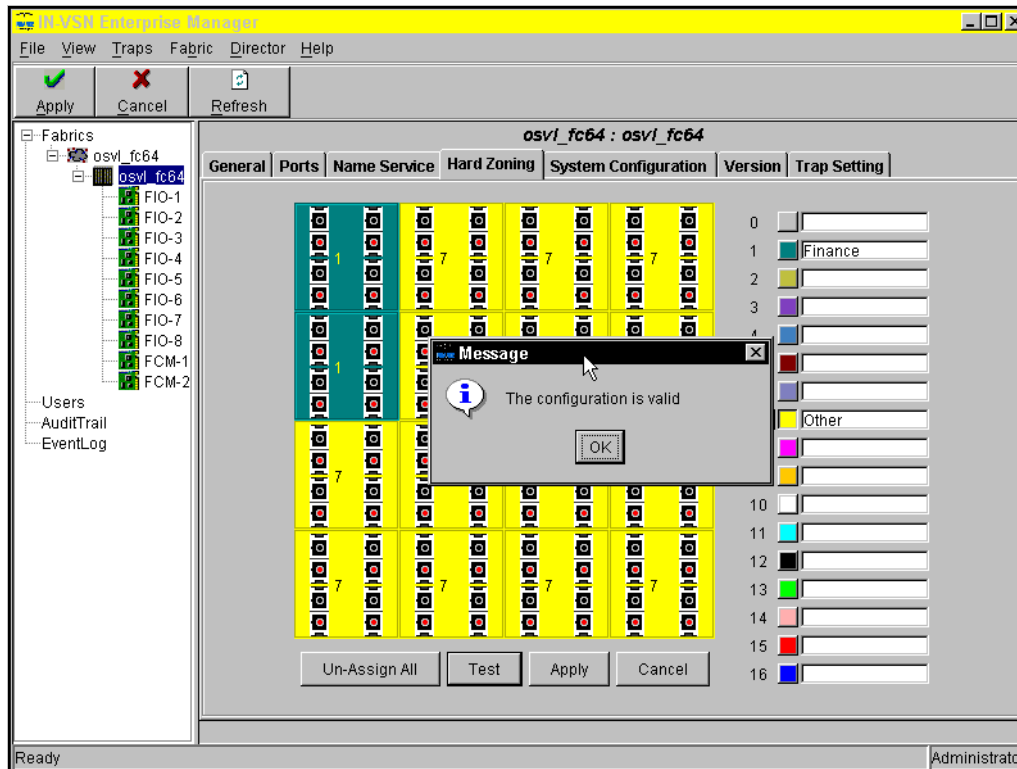


Figure 3-63 IN-VSN: Having two hard zones defined

We now have two zones defined: *Finance* and *Other*.

The *Finance* zone consists of two port groups giving a total port count of eight. The *Other* zone has the remaining 14 port groups assigned to it having a total port count of 56.

However, this hard zone setup is not yet active since we have not clicked the **Apply** button at this stage.

To demonstrate the adjoining rule, we have added the bottom left port group to the *Finance* hard zone.

Remember that the adjoining rule for hard zoning demands that all port groups of one zone must be adjoining.

As illustrated in Figure 3-64, when clicking the **Test** button now we get a violation message, as expected.

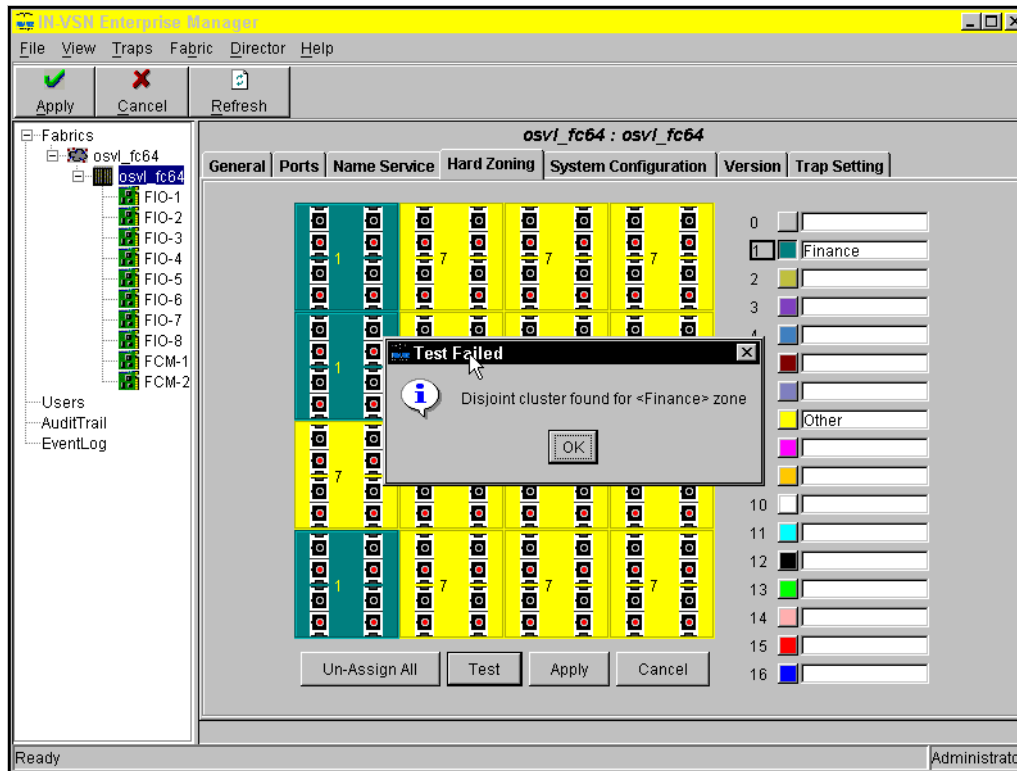


Figure 3-64 Violation of ports

To successfully finish our hard zone setup as planned, we have put the bottom left port group back into the *Other* zone.

To activate our hard zone layout, we click **Apply** as illustrated in Figure 3-65.

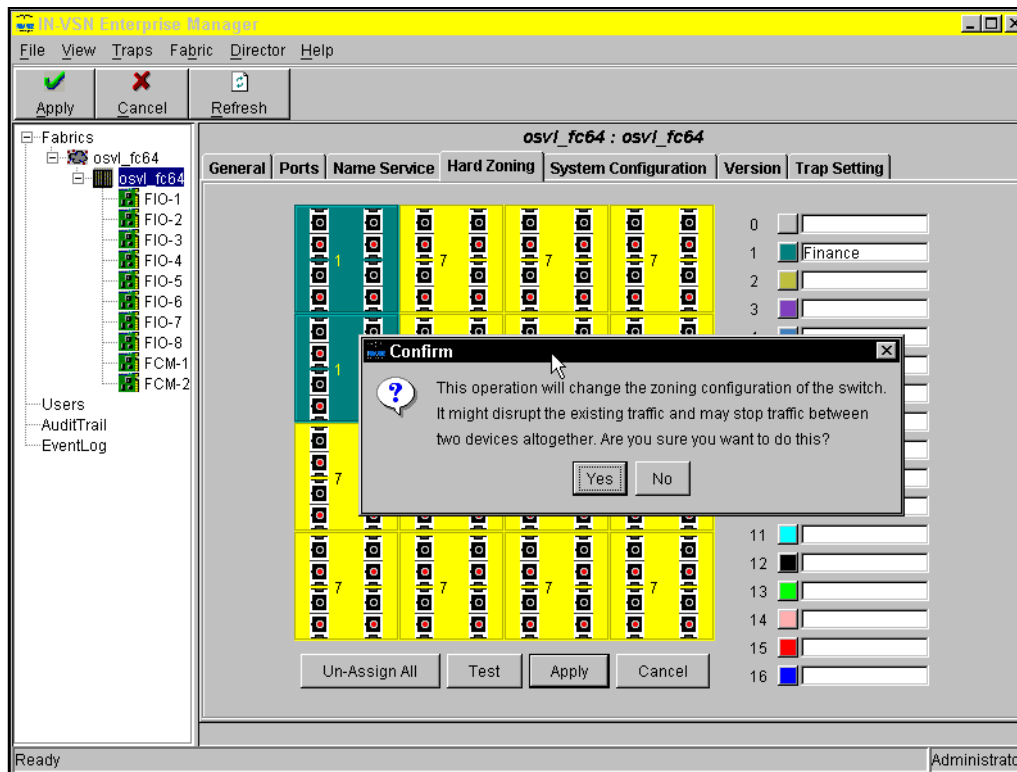


Figure 3-65 IN-VSN: Applying a hard zone setup

As explained earlier, a director reconfiguration task such as this will affect in-progress I/O. A warning message is issued to indicate this.

To activate our setup, we click **Yes**.

After this hard zone setup is activated, there is no communication possible between the *Finance* zone and the *Other* zone.

This is the most secure way to separate nodes in a fabric environment.

3.4.6 Defining name server zones

Our goal is to enable all four HP servers to access the ESS.

To achieve a high level of availability, we have chosen to implement a multipathing architecture except for one server (GEODE).

All *primary* FC ports of these servers will be zoned to ESS Bay 1, Port 4, and all *secondary* FC ports of these servers will be zoned to ESS Bay 4, Port 4.

We use **pvlinks** as multipathing software. Another possibility is to use the IBM Subsystem Device Driver (SDD).

In Figure 3-66 we show an overview of what we will achieve.

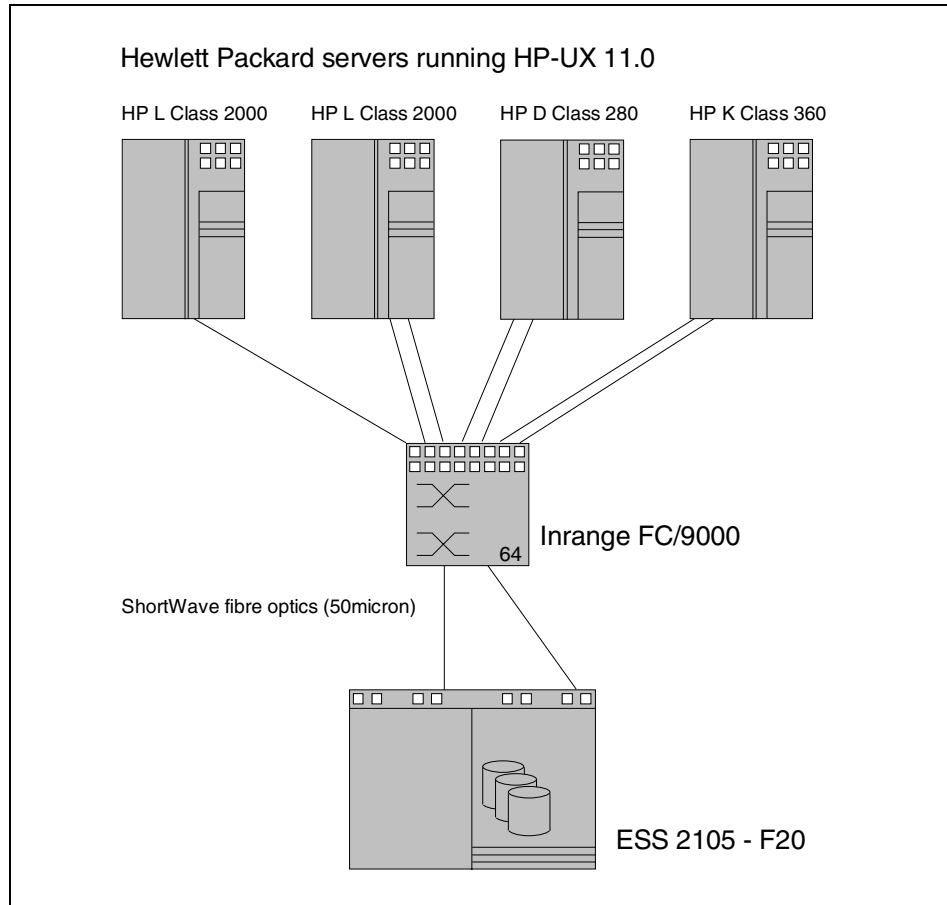


Figure 3-66 Server setup for soft zoning

Be aware that such a zone-configuration is only advisable in specific environments.

Be sure to have only host adapters of the same type in a zone. This is true in our case. Therefore, we just implement two zones instead of dedicated zones for each server.

By implementing two zones, we limit the access of a particular host FC port to only one ESS FC port. We do this because we like to have only two paths to a ESS logical volume.

These two paths are created by giving the two host FC ports of one server access to the particular ESS logical volume using ESS LUN masking.

In Figure 3-67 we show how these ports are physically attached to the INRANGE director.

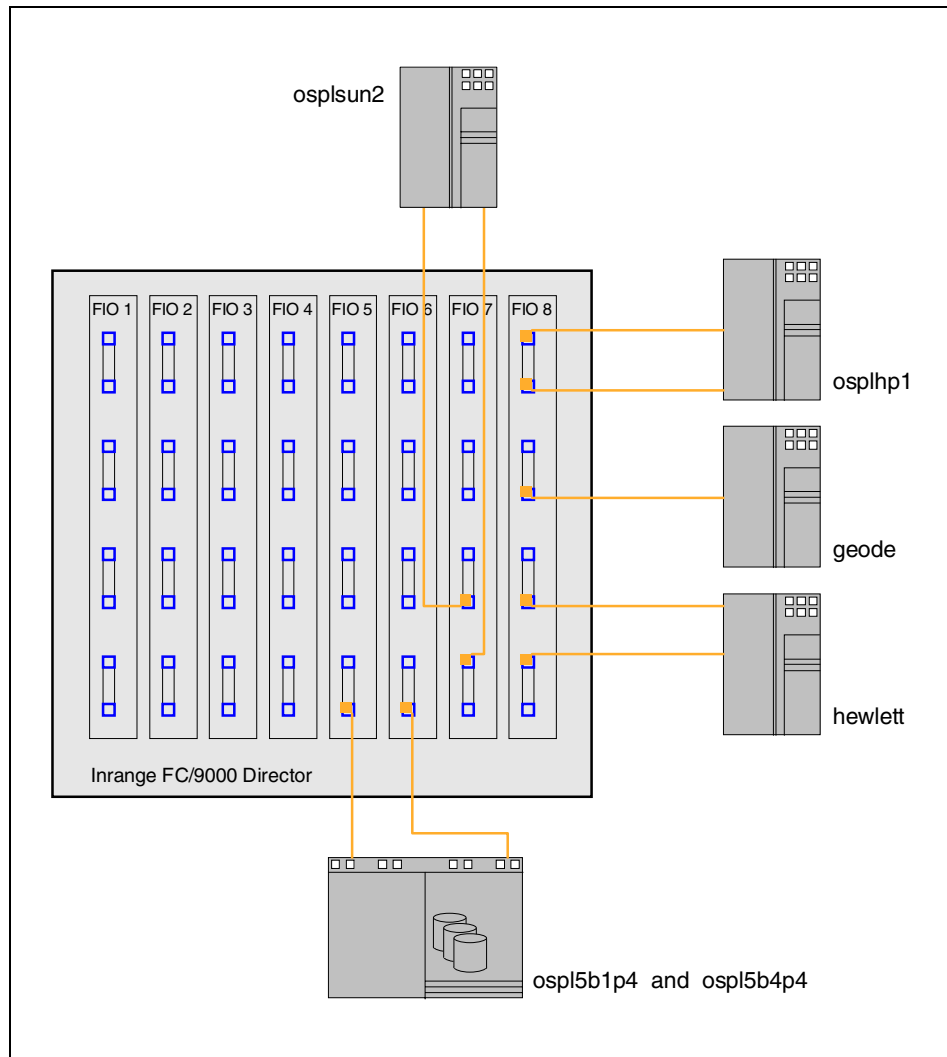


Figure 3-67 Physical cable connection for soft zoning setup

We shall now explain how our zoning will look logically. As illustrated in Figure 3-68, we plan to have two zones.




Figure 3-68 Logical view of our two name server zones

Zone *ospl5_b1_zone* has four director ports as members, including the one ESS port b1p4 is attached to.

Zone *ospl5_b4_zone* has 5 director ports as members, including the port that is used to attach ESS port b4p4.

By implementing such a layout, each server is attached to two zones. Since each zone has a different ESS port as a member, the server will get exactly two paths to the ESS. Use of **pvlinks** and ESS based LUN masking will enable sufficient multipathing.

To actually implement name server zoning, click the specific fabric icon, in our case *osvl_fc64*, in the navigation tree, and then select the **Zoning** tab as shown in **Figure 3-69**.

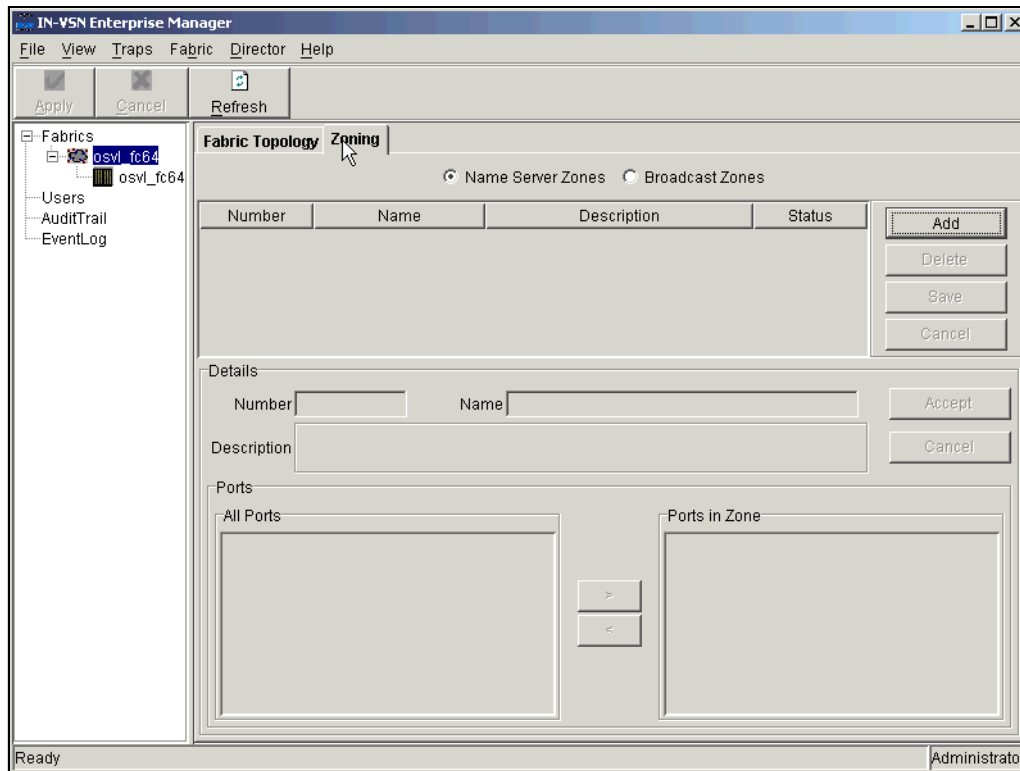


Figure 3-69 IN-VSN: Entering the zone screen in fabric view mode

Initially, there are no zones defined at all. Remember that in such a case, the only effective zone is the default zone. This enables all ports to communicate with all other ports (any-to-any). Only zoned ports will have limited but controlled access. All unzoned ports are considered as being part of the default orphan zone, which enables them to communicate without any restrictions.

Leave the **Name Server Zone** field selected. Click **Add** to create a new zone. Enter zone number, zone name, and a description as illustrated in Figure 3-70.

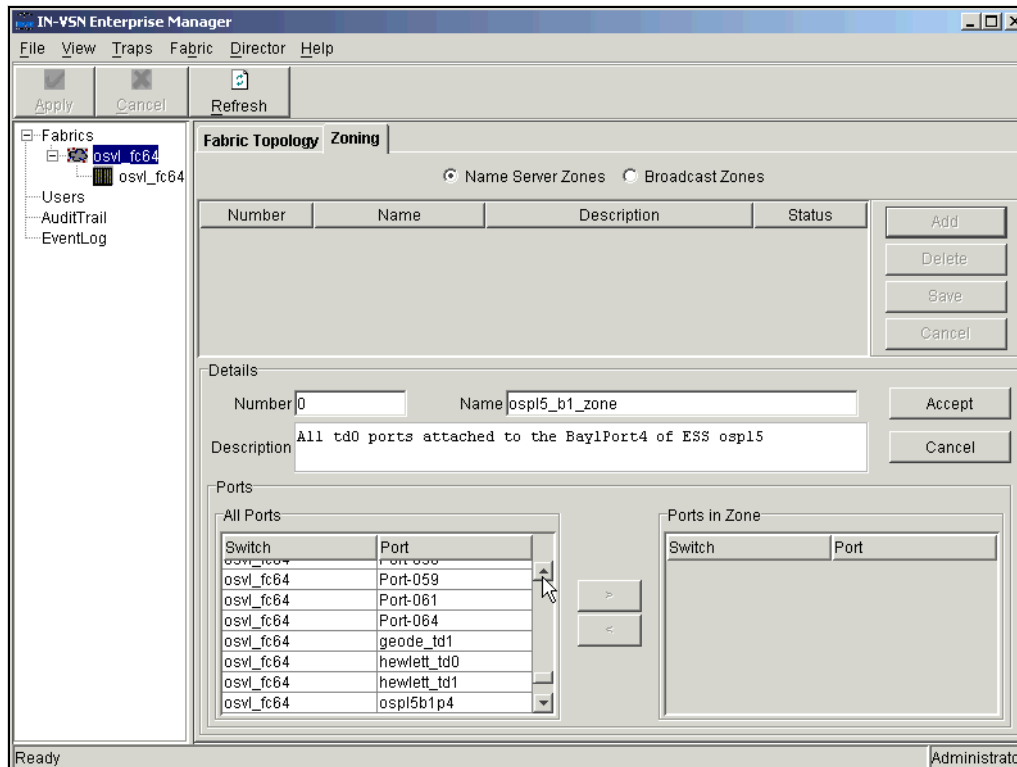


Figure 3-70 IN-VSN: Entering number and name for a new zone

From the Ports scroll list, select the ones you want to be a member of that zone as shown in Figure 3-71.

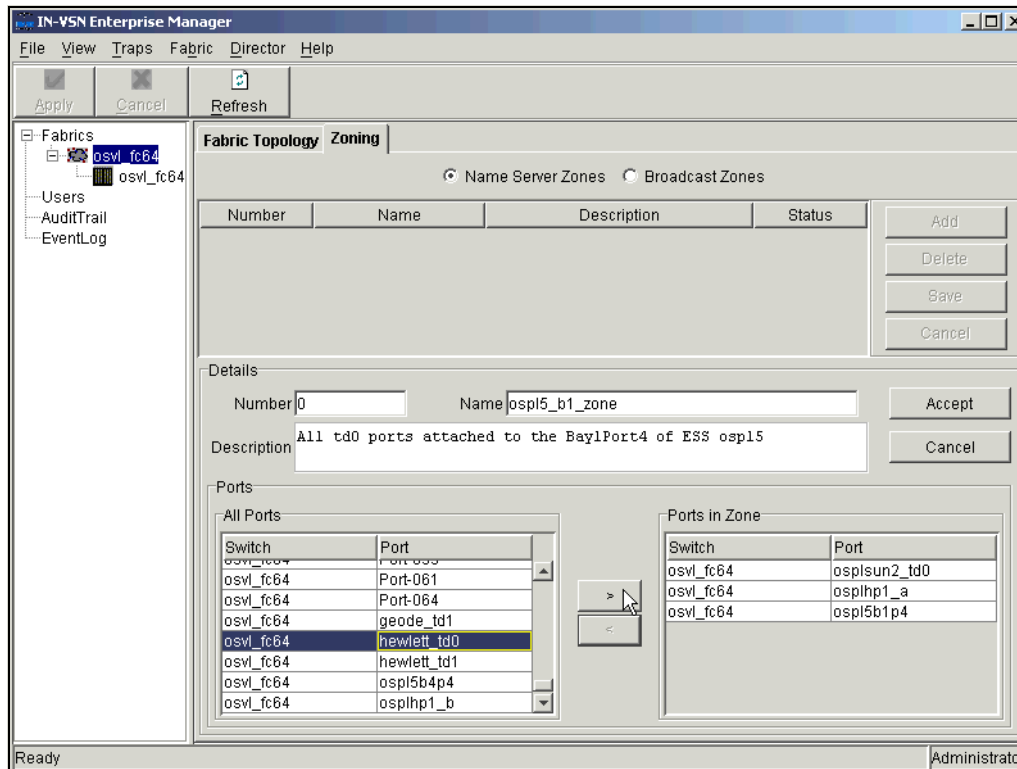


Figure 3-71 IN-VSN: Selecting the members of a new zone

We added these ports to the zone *ospl5_b1_zone*:

- ▶ osplsun2_tdo
- ▶ osplhp1_a
- ▶ ospl5b1p4
- ▶ hewlett_tdo

The effect is that all these server ports will have access only to the ESS port Bay1Port4.

After putting all needed ports in the **Ports in Zone** list, we continue with clicking **Accept**. This is shown in Figure 3-72.

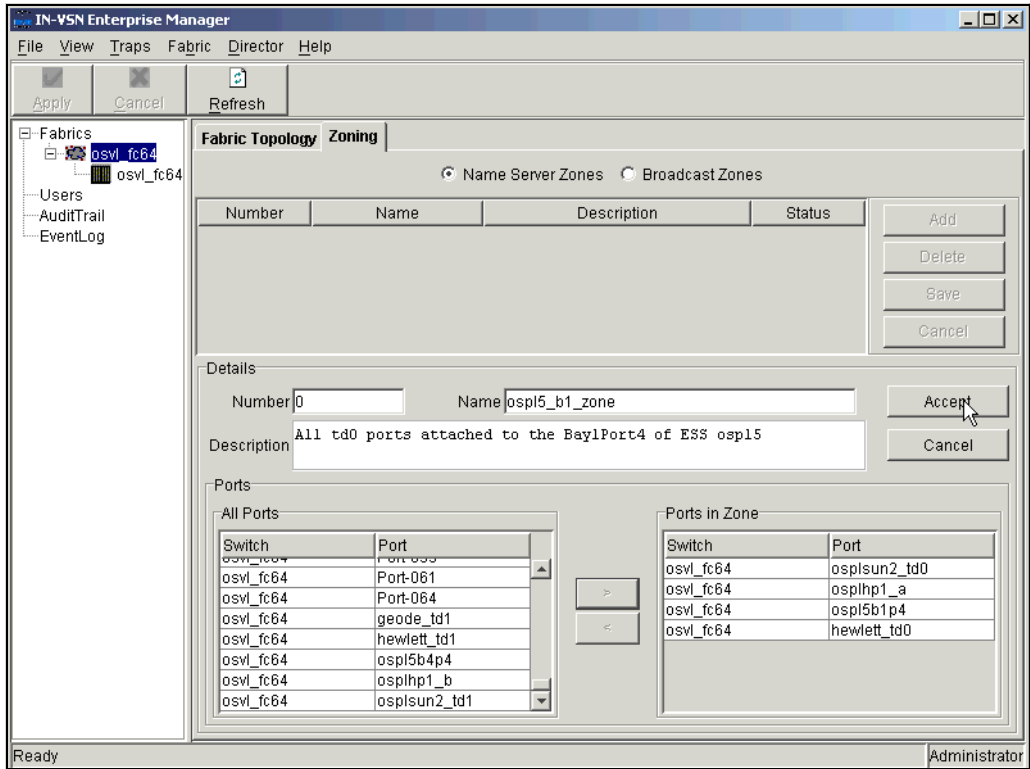


Figure 3-72 IN-VSN: Accepting settings for first new zone

However, now we are getting an error message, since our previously chosen zone number of **0** is not valid. This is shown in Figure 3-73.

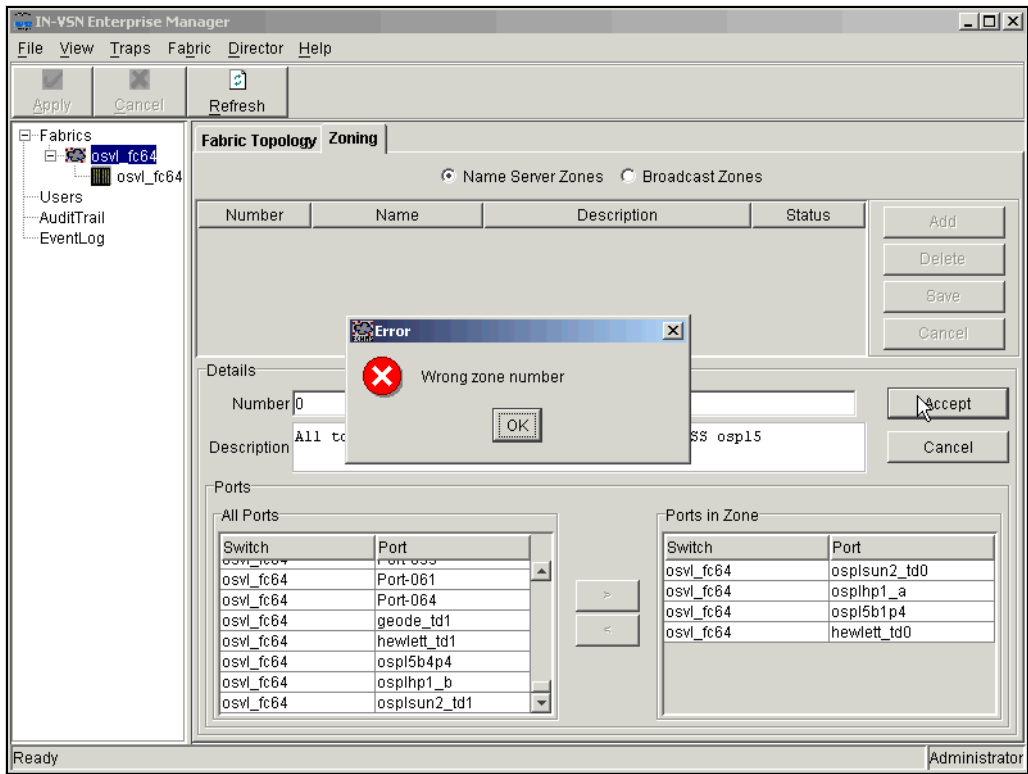


Figure 3-73 IN-VSN: Zone number 0 is not allowed to use

So we just change this zone number to **1**, and click **Accept** again. Now this action is accepted.

This zone is now listed in the zone list with a status of *Added*.

To make this zone effective, click **Save** as shown in Figure 3-74.

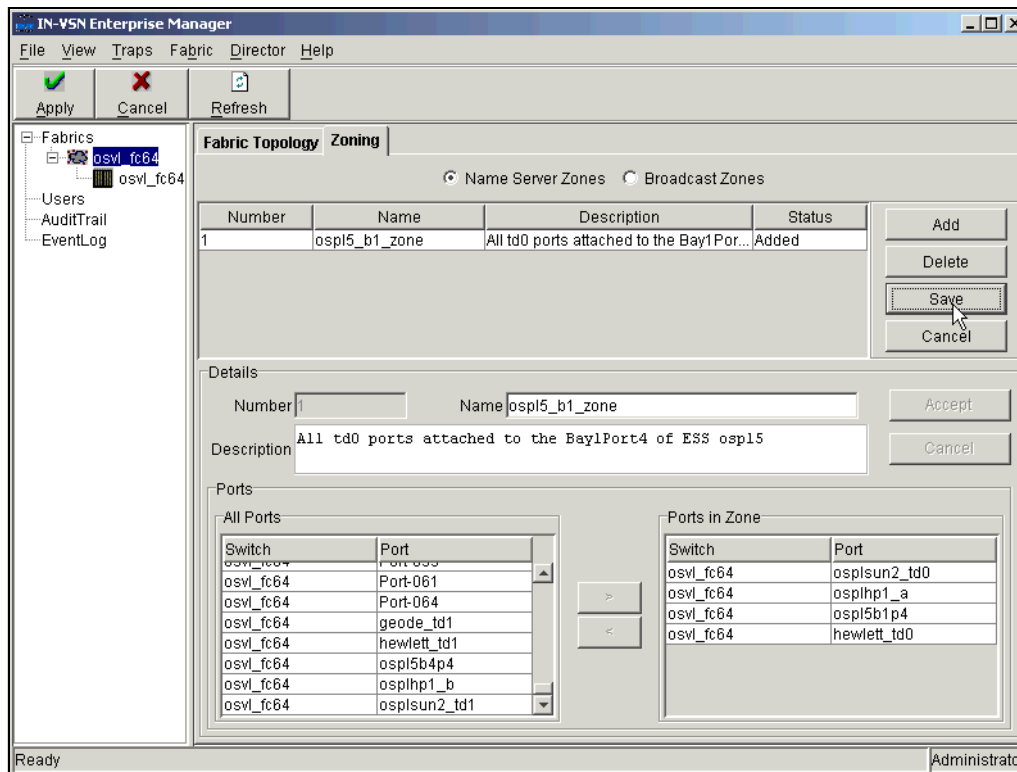


Figure 3-74 IN-VSN: Saving a newly added zone

Using the same procedure, we now add a second zone called *osp15b4_zone* which has these ports as members:

- ▶ osp15b4p1
- ▶ osp1sun2_td1
- ▶ hewlett_td1
- ▶ geode_td1
- ▶ osp1hp1_b

After saving this second zone, the zone list contains both zones with a status of Saved as illustrated in Figure 3-75.

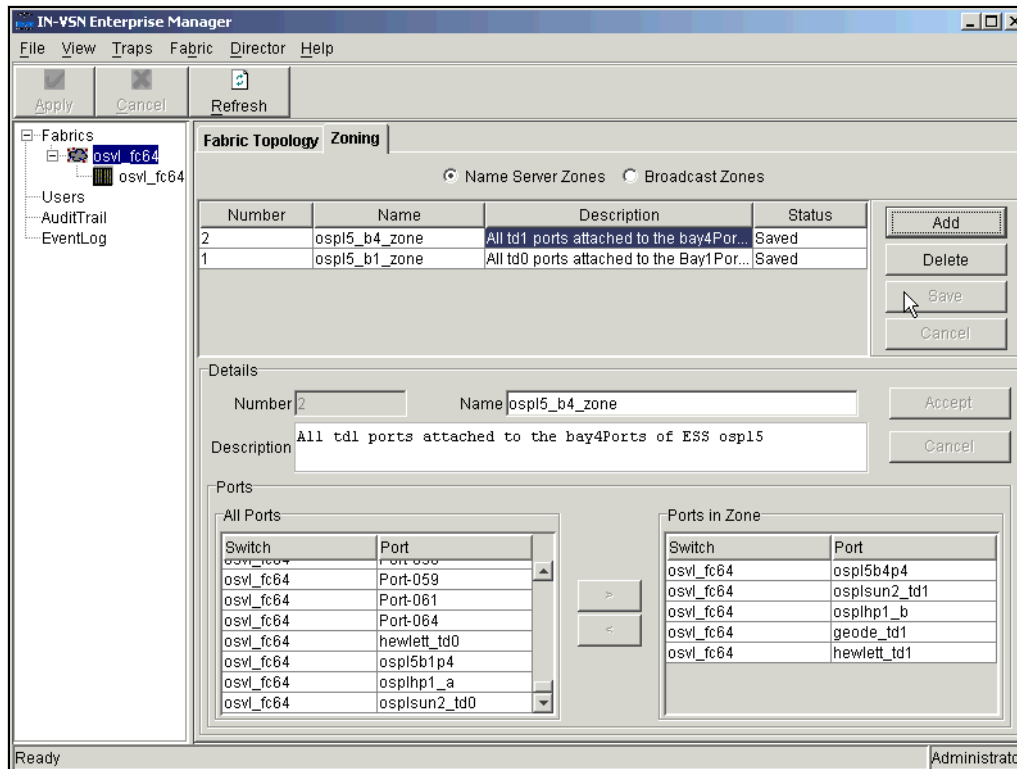


Figure 3-75 IN-VSN: Zone list with both zones added and saved

Another way to verify the right zone settings is to have a look at the Ports View of the director.

To do this, select the specific director in the navigation tree and then click the **Ports** tab. After scrolling down we see the used ports with our given names, the WWN of the attached ports, and then the zone these ports belong to as shown in Figure 3-76.

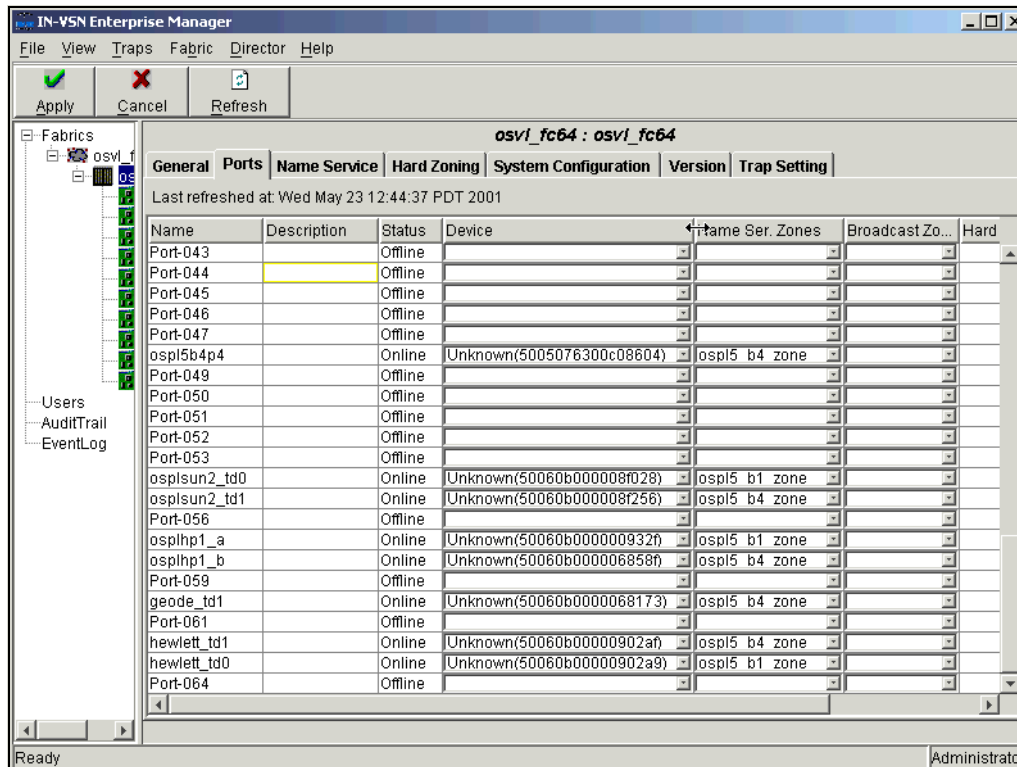


Figure 3-76 IN-VSN: Director Port View showing the port to zone relationship

We now have two zones defined in a fabric consisting of one director.

Since both zones are saved and all specified external ports are actually attached and online, we can start I/O between these servers and our ESS.

3.5 Cascading of directors

To further increase the amount of available FC ports, we can connect multiple INRANGE directors using FC links. This kind of FC link will be automatically discovered. This link is widely called an Inter-Switch Link (ISL).

The process of connecting directors using ISLs is known as cascading.

The motivation to use cascaded directors is not limited to increased port counts. Cascaded directors can be used to create a SAN that spans large distances.

Unique to INRANGE Directors, ports used for Inter Switch Links are called T_Ports (Trunk ports). Every director port can be used for T_Port operation. No manual setting is needed since ISL connections are automatically detected and the affected ports are set to T_Port.

ISLs can be run on either ShortWave GBICs or LongWave GBICs. The usage of LongWave GBICs enables the creation of long distance fabrics. The actual supported distance between INRANGE directors can be further expanded by using channel extenders or multiplexers.

Figure 3-77 shows an example where cascaded directors are used to allow host based mirroring over long distances:




Figure 3-77 One INRANGE fabric consisting of two cascaded directors

Cascading is a very effective and easy to use method to extend distances and increase maximum port count. However we need to consider some of the implications which are:

- By connecting multiple directors they will become one fabric inheriting all the rules for one fabric, including:
 - Name server zones are valid fabric-wide.
 - Numbers and names of name server zones must be unique fabric-wide.
 - Either all or none of the director ports in a fabric are hard zoned.

- Each director or switch in a fabric must have a unique switch id.
 - When having no zoning implemented at all then all nodes have access to any other node fabric-wide.
- ▶ When actually using Inter Switch Links for I/O, there will be a delay called latency. The actual usage of a ISL is commonly called a hop. Typically up to three hops are supported for a particular I/O.
 - ▶ T_Ports cannot be part of name server zones.
 - ▶ When having hard zones installed that cross chassis boundaries, then both parts of the hard zone must be directly connected by at least one ISL. You must have dedicated ISLs for each hard zone spanning multiple chassis.
 - ▶ Up to eight FC/9000 directors can be cascaded to create one fabric.

All directors and switches being part of one fabric will be displayed with their ISLs when having the **IN-VSN** fabric view selected, as illustrated in Figure 3-78.

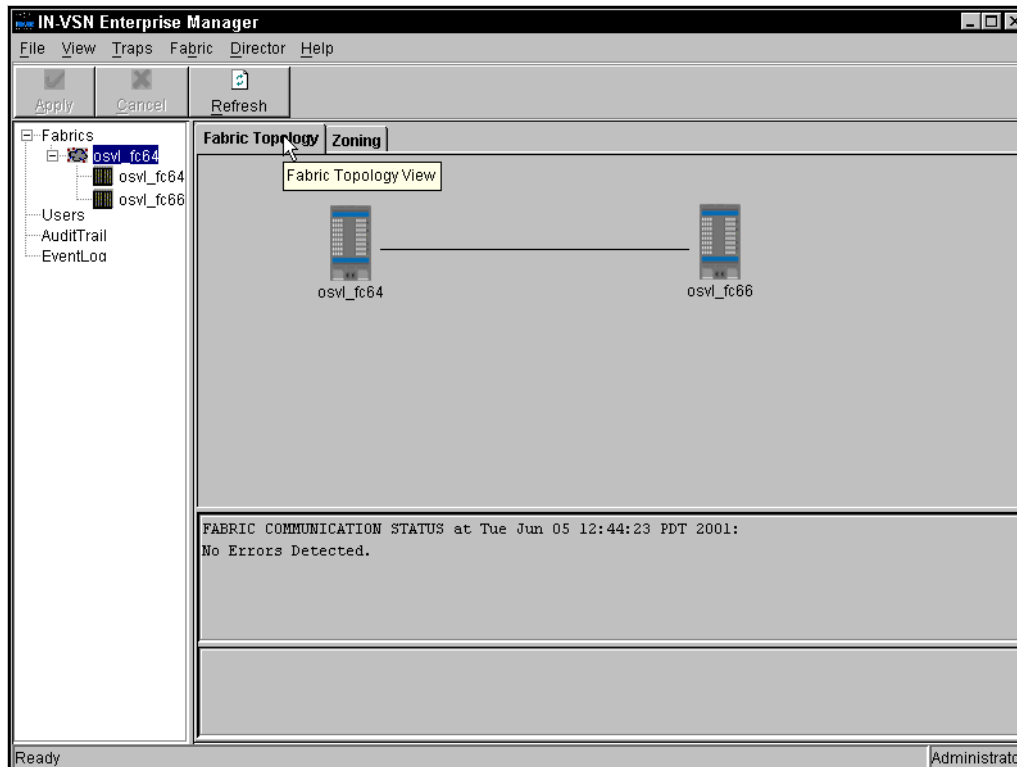


Figure 3-78 IN-VSN: Fabric view with two cascaded directors

3.6 Monitoring and maintenance of an INRANGE SAN

This topic will briefly cover these topics:

- ▶ Communication protocols to manage an INRANGE SAN
- ▶ Microcode load information
- ▶ Monitoring user activities
- ▶ Using the INRANGE event log

3.6.1 Management communication protocols

Management communication can be done using different physical interfaces:

- ▶ Director:
 - RS-232 Serial Interface
 - Ethernet 10/100 Base T Connector
- ▶ Management PC:
 - Modem connectivity
 - RS-232 Serial Interface
 - Ethernet 10/100 Base T Connector to directors
 - Industry standard network interfaces to the corporate network

Using these physical interfaces the following management services can be used:

IN-VSN management software

Client/Server based management tool dedicated for users.

Maintenance interface using the RS-232 serial interface

Terminal based interface to change basic settings like director ID's or IP settings. This is dedicated for IBM or INRANGE Customer Engineers.

Simple network management protocol (SNMP)

An interface to integrate INRANGE directors into the corporate systems management.

Trivial file transfer protocol (TFTP)

TFTP is used to load new firmware to directors and to retrieve and/or change director management configuration settings (for example, IP, SNMP). This service is dedicated to IBM or INRANGE Customer Engineers.

Call home functionality

Call home enables the IN-VSN software to propagate event information and configurations to INRANGE or IBM systems. This allows better response times in case of component failures. Call home is disabled by default but can be enabled by Customer Engineers. A modem connection for the Management PC is required.

Changing the IP address

The IP settings of INRANGE directors should be done by IBM or INRANGE Customer Engineers.

The change of the IP settings of an INRANGE FC/9000 director is performed using a Terminal session via the RS232 Serial Interface.

Each FCM blade has its own IP settings and also its own Serial interface.

You should always keep both IP settings valid to ensure management access in the event of FCM failures.

To change both addresses, you have to login to both FCM blades one after the other. So, first you login to one FCM to set its IP settings and then you login to the other FCM to change its IP settings.

The change process for IP settings involves a booting of the FCM. However this local FCM boot process does not affect other director parts. Neither the alternate FCM blade nor the director functions will be disturbed.

This results in a non-disruptive IP change.

3.6.2 Microcode-loads

Microcode-loads can be done non-disruptively using the IN-VSN Enterprise Manager in conjunction with the TFTP interface. However, this action should be performed by IBM or INRANGE customer engineers.

3.6.3 Monitoring user activities

IN-VSN management software allows to monitor user activities using an Audit Trail.

All user levels (*admin*, *oper*, *viewer*) can access the audit trail.

All activities are logged and categorized into different types. Following are some examples of these operation types:

- ▶ User login
- ▶ IP address of user login
- ▶ Fabric definition
- ▶ User definitions
- ▶ Name server zoning
- ▶ Switch name changed

To look at the audit trail, click **AuditTrail** in the navigation tree as illustrated in Figure 3-79.

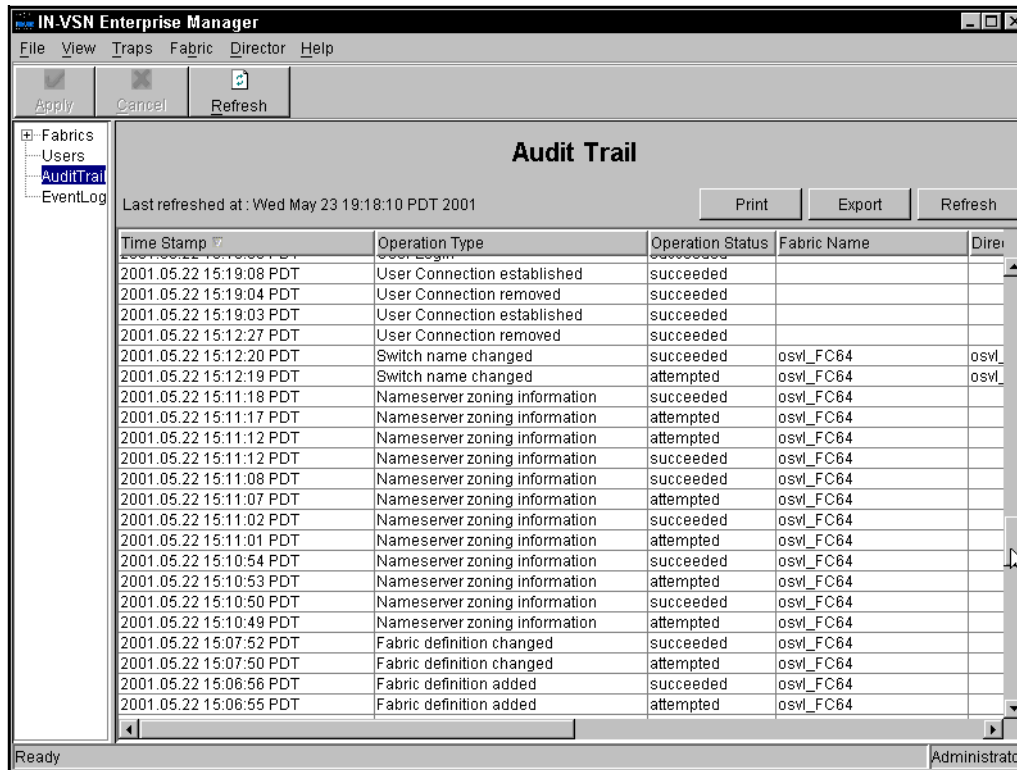


Figure 3-79 IN-VSN: Using the INRANGE Audit-Trail to monitor user activities

3.6.4 Using the IN-VSN event log

The event log contains all important events that have occurred. This includes events triggered by users and events caused by other external or internal influences, such as FRU failures or losing power.

To read the event log, click **EventLog** in the navigation tree of IN-VSN, as shown in Figure 3-80.

Event ID	Time	Fabric Name	Director Name	Event Description	FRU	S
625	2001.06.04 06:15:58 P...	osvl_fc64	osvl_fc64	TIME SYNC OCCURRED(ACM)	FCM1	5
624	2001.06.04 05:40:04 P...	osvl_fc64	osvl_fc64	TIME SYNC OCCURRED(ACM)	FCM1	5
623	2001.06.04 00:15:58 P...	osvl_fc64	osvl_fc64	TIME SYNC OCCURRED(ACM)	FCM1	5
622	2001.06.03 23:40:04 P...	osvl_fc64	osvl_fc64	TIME SYNC OCCURRED(ACM)	FCM1	5
44	2001.06.03 21:40:11 P...	N/A	N/A	FRU call home failure. FRU Call H...		4
621	2001.06.03 18:15:58 P...	osvl_fc64	osvl_fc64	TIME SYNC OCCURRED(ACM)	FCM1	5
620	2001.06.03 17:40:04 P...	osvl_fc64	osvl_fc64	TIME SYNC OCCURRED(ACM)	FCM1	5
619	2001.06.03 12:15:58 P...	osvl_fc64	osvl_fc64	TIME SYNC OCCURRED(ACM)	FCM1	5
618	2001.06.03 11:40:04 P...	osvl_fc64	osvl_fc64	TIME SYNC OCCURRED(ACM)	FCM1	5
617	2001.06.03 06:15:58 P...	osvl_fc64	osvl_fc64	TIME SYNC OCCURRED(ACM)	FCM1	5
616	2001.06.03 05:40:04 P...	osvl_fc64	osvl_fc64	TIME SYNC OCCURRED(ACM)	FCM1	5
615	2001.06.03 00:15:58 P...	osvl_fc64	osvl_fc64	TIME SYNC OCCURRED(ACM)	FCM1	5
614	2001.06.02 23:40:04 P...	osvl_fc64	osvl_fc64	TIME SYNC OCCURRED(ACM)	FCM1	5
43	2001.06.02 21:40:11 P...	N/A	N/A	FRU call home failure. FRU Call H...		4
613	2001.06.02 18:15:58 P...	osvl_fc64	osvl_fc64	TIME SYNC OCCURRED(ACM)	FCM1	5
612	2001.06.02 17:40:04 P...	osvl_fc64	osvl_fc64	TIME SYNC OCCURRED(ACM)	FCM1	5
611	2001.06.02 12:15:58 P...	osvl_fc64	osvl_fc64	TIME SYNC OCCURRED(ACM)	FCM1	5
610	2001.06.02 11:40:04 P...	osvl_fc64	osvl_fc64	TIME SYNC OCCURRED(ACM)	FCM1	5
609	2001.06.02 06:15:58 P...	osvl_fc64	osvl_fc64	TIME SYNC OCCURRED(ACM)	FCM1	5
608	2001.06.02 05:40:04 P...	osvl_fc64	osvl_fc64	TIME SYNC OCCURRED(ACM)	FCM1	5
607	2001.06.02 00:15:58 P...	osvl_fc64	osvl_fc64	TIME SYNC OCCURRED(ACM)	FCM1	5
606	2001.06.01 23:40:04 P...	osvl_fc64	osvl_fc64	TIME SYNC OCCURRED(ACM)	FCM1	5
42	2001.06.01 21:40:11 P...	N/A	N/A	FRU call home failure. FRU Call H...		4

Figure 3-80 IN-VSN: Accessing the INRANGE event log

This event log is valid for all fabrics managed by this IN-VSN server. It includes the logs of multiple fabric devices and even multiple fabrics.

To file this log for future usage, you are able to export this log. Two file types can be used to export the log:

- ▶ Comma Separated Value Files (*.csv)
- ▶ Text Files (*.txt)

3.7 ESS configuration guidelines

In this chapter we have addressed some of more commonly encountered features of the INRANGE director.

The 2042-001 is designed to provide maximum bandwidth switching from any port in the system to any other port in the system. There is no mandatory requirement for attachment of server or device ports to the 2042-001.

There are, however, a few guidelines to system configuration that should be followed to assure optimum system reliability, availability and serviceability of the overall SAN environment to which the 2042-001 is installed.

The foremost consideration is to spread critical resources across multiple Director FRUs. The 2042-001 FRU in this instance is the eight port FI/O module. Device ports should be distributed over as many director FI/Os as available. Assignment of device ports to the FI/O is not critical and assignment to the same port number across adjacent FI/Os for convenience is acceptable. When device ports are allocated to the director, allocation of additional ports in the same horizontal row should be considered for future expansion.

Channels from a common server or host image should be allocated over multiple FI/Os in a manner similar to the way device ports are distributed over multiple FI/Os as described above. In addition, when a server channel is daisy chained to multiple devices by IOCP, the channel and device ports may be grouped on an FI/O.

Critical system resources, such as SYSRES storage, may be assigned as device and channel pairs to take advantage of port pairs on a FI/O that offer the greatest autonomy of routing. These port pairs are ports 1&2, 3&4, 5&6, 7&8 on every FI/O. Selection of these ports offers the highest availability in the most adverse conditions.



Implementing the SAN Data Gateway

In this chapter, we describe the steps involved in planning and implementing the IBM Storage Area Network Data Gateway, 2108-G07. The SAN Data Gateway is a hardware solution to allow connection of Fibre Channel ready host systems to attach to SCSI storage systems.

4.1 SAN Data Gateway

The IBM Storage Area Network Data Gateway is an essential component of the SAN infrastructure. It provides several benefits to bridge the legacy gap as storage products migrate from SCSI based attachments to Fibre Channel.

A diagram to show a SAN Data Gateway configuration using a single host is shown in Figure 4-1.




Figure 4-1 SAN Data Gateway configuration

The IBM Storage Area Network Data Gateway allows you to:

- Protect legacy storage equipment while utilizing the latest host servers with Fibre Channel support

- ▶ Expand connectivity to storage devices with use of IBM SAN hubs, switches, and directors
- ▶ Perform channel zoning and LUN masking capability to allow access at a volume level
- ▶ Overcome the distance limitations of SCSI based host systems using longwave ports that support distances up to 10 km
- ▶ Utilize the StorWatch SAN Data Gateway Specialist which is an easy to use interface for managing and controlling access of host systems to storage devices

The SAN Data Gateway is available as a rack-mount unit or as a stand-alone tabletop unit. The gateway model provides two shortwave Fibre Channel ports and four Ultra SCSI Differential ports to attach disk or tape storage devices. One or two Fibre Channel cards — dual-port, shortwave and/or single port, longwave — may be added for a maximum of six shortwave ports, or two shortwave and two longwave ports. If you are using the dual-port shortwave cards, Figure 4-2 depicts the port assignment numbers for the optical interfaces.




Figure 4-2 SAN connection port assignment

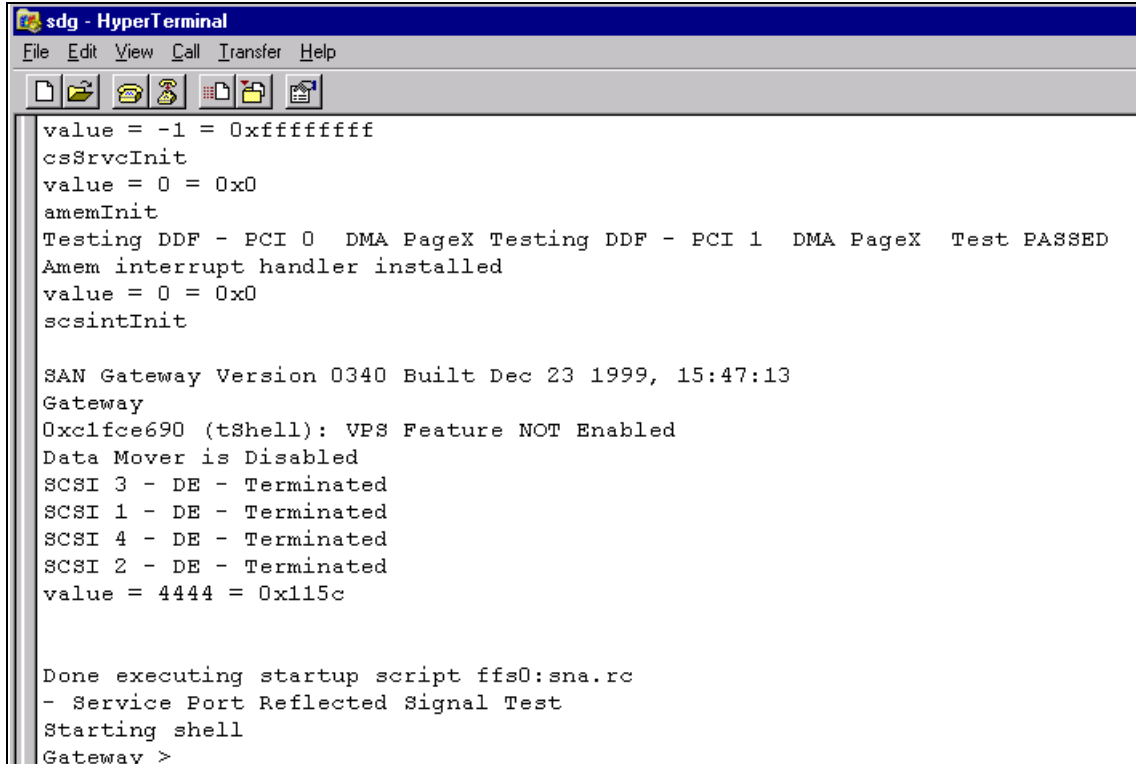
4.2 Installation

Before any server or device connection is made to the SAN Data Gateway, you need to power on the unit and connect a service terminal to the 9-pin Service port located on the rear right-hand side of the unit. A PC running a terminal emulation program, such as Windows HyperTerminal or NETTERM, can be used. The settings of the Service port are:

- ▶ 19200 baud
- ▶ 8 data bits
- ▶ No parity
- ▶ 1 stop bit
- ▶ X-on/X-off flow control
- ▶ VT-100 compatible terminal

If a PC with terminal emulation is used, a 9-pin female to 9-pin female, null modem cable is required and is provided with the unit.

Once connected, power on the SAN Data Gateway and the start up messages will appear and scroll across the window. When the power on sequence has completed, a prompt **Gateway>** appears on the window, as shown in Figure 4-3.



```
sdg - HyperTerminal
File Edit View Call Transfer Help

value = -1 = 0xffffffff
csSrvInit
value = 0 = 0x0
amemInit
Testing DDF - PCI 0 DMA PageX Testing DDF - PCI 1 DMA PageX Test PASSED
Amem interrupt handler installed
value = 0 = 0x0
scsintInit

SAN Gateway Version 0340 Built Dec 23 1999, 15:47:13
Gateway
0xc1fce690 (tShell): VPS Feature NOT Enabled
Data Mover is Disabled
SCSI 3 - DE - Terminated
SCSI 1 - DE - Terminated
SCSI 4 - DE - Terminated
SCSI 2 - DE - Terminated
value = 4444 = 0x115c

Done executing startup script ffs0:sna.rc
- Service Port Reflected Signal Test
Starting shell
Gateway >
```

Figure 4-3 IBM Storage Area Network Data Gateway startup

If you type in **help** and then press Enter, a list of available commands is provided. The commands are case sensitive and *must* be entered as they appear.

Issue the **initializeBox** command to remove any configurations files that may be present. The unit will restart automatically.

Note: The **initializeBox** command will erase all configuration files. It should be used only during initial SAN Data Gateway installation.

4.2.1 Setting the Ethernet address

Once restarted, the Ethernet port must be configured and attached using network information provided by the network administrator. To set the IP address, use the **ethAddrSet** command. The address must contain the double quotes (“):

```
Gateway > ethAddrSet "9.111.24.66"  
Network not Enabled  
Write complete  
Host Address set to 9.111.24.66 for Ethernet interface  
value = 0 = 0x0  
Gateway >
```

If a subnet is required then add it after the IP address and separate the two addresses using a comma. For example:

```
Gateway > ethAddrSet "9.111.24.66", "255.255.255.0"
```

If a gateway or standard router is to be specified, then issue the **gateAddrGet** command to view if there is one set and the **gateAddrSet** command to set or change it. For example:

```
Gateway > gateAddrGet  
No current gateway address set  
value = 0 = 0x0  
Gateway > gateAddrSet "193.1.1.11"  
Write complete
```

The Ethernet port on the SAN Data Gateway comes from the factory disabled. To enable it, you must issue the **ethEnable** command. This will not take effect until the unit is rebooted. The reboot can occur from a power off, or by issuing the **reboot** command. During the reboot, you will see that the IP address is set and now enabled.

```
Gateway > ethEnable  
Write complete  
Ethernet will be enabled on next boot  
value = 0 = 0x0  
Gateway > _
```

4.2.2 Using Telnet on the SAN Data Gateway

If a user would prefer to telnet to the SAN Data Gateway rather than by using the service terminal port after initial setup, this can be done. First you must create a user from the service terminal by using the **userAdd** command. Enter the login name and password using the quotes and comma:

```
ITS0 > userAdd "itso", "residency"  
value = 0 = 0x0  
ITS0 > _
```

You cannot telnet to the Gateway and use the service port at the same time. When you telnet to the Gateway, the service port on the rear of the unit will stop its communications. After you end the telnet session, then the service port will become available again.

4.2.3 Startup sequence

You must start up the SAN Data Gateway and the attached host and target devices in a specific order. When you add or remove SCSI devices or update firmware, you must restart. The following procedures describe the situations and order of procedure when you restart the SAN Data Gateway.

Before you restart the SAN Data Gateway, you must stop all input and output (I/O) activity between the host and SCSI devices.

1. SCSI devices

Turn on the SCSI devices. You must turn on all SCSI devices attached to the SAN Data Gateway before you initially turn on or restart the SAN Data Gateway.

2. SAN Data Gateway

The SAN Data Gateway scans the SCSI buses when it starts. If you add or remove SCSI devices after the Gateway has started, the Gateway will not detect the changes. You can invoke an SCSI rescan or restart operation from either the StorWatch SAN Data Gateway Specialist client or the service terminal.

3. Fibre Channel host

Before you turn on or restart the hosts that are connected with Fibre Channel to the SAN Data Gateway, you must wait until the SAN Data Gateway has finished starting. You will know the Gateway has finished starting when the ready light on the front panel blinks at frequency intervals of one second.

- Some operating systems provide you with software methods that allow you to add or remove SCSI devices dynamically after the host has started. To ensure reliable operation, restart the host.
- If you update SAN Data Gateway firmware, you must restart the Gateway to use the new firmware. To ensure compatibility between the firmware features or functions and the host, restart the host.
- If you update SCSI device firmware, the SAN Data Gateway Explorer application does not display the new firmware version until the SAN Data Gateway has issued an SCSI inquiry. The SCSI inquiry occurs when the Gateway rescans the SCSI buses. The SCSI inquiry also occurs when the StorWatch SAN Data Gateway Specialist client application or the service terminal rescans the SCSI buses.

Currently, up to eight different hosts can be attached to each Fibre Channel port. If all six ports are installed, then 48 different hosts can attach to the SAN Data Gateway.

4.3 StorWatch SAN Data Gateway Specialist

The StorWatch SAN Data Gateway Specialist software provides remote capability for all management, configuration, and event notification. It is comprised of three parts:

- ▶ Agent
- ▶ Server
- ▶ Client

Agent

The agent is embedded in the operating system of each SAN Data Gateway to provide a stand-alone manageable host. The StorWatch SAN Data Gateway Specialist software uses SNMP to set and retrieve information that controls the operation of the Agent. The Specialist also uses SCSI over TCP to allow updates to the Gateway and target device.

Server

The server is a Java application that runs on a host and is used to maintain communication with the agents and acts as an intermediary between the agent and the client. The server coordinates the request from multiple clients to manage multiple gateways or agents. Multiple clients can share data the server already knows about, and the server receives all traps from the agent and forwards them to the clients that are registered to receive them.

Client

The client is a Java application that operates from any compatible computer as long as it has a TCP/IP connection to the server. One or more clients can connect to a server. The client provides the user interface to allow the management and configuration of the SAN Data Gateway.

The server and client can be installed on to the same computer.

The StorWatch SAN Data Gateway Specialist supports the following operating systems:

- ▶ Windows 95,98, 2000, and NT 4.0 with SP5 or later
- ▶ AIX ver 4.3.3 or later
- ▶ Solaris 2.6 or later

4.3.1 Installing StorWatch Specialist

The Specialist software is not bundled with the SAN Data Gateway. The Specialist software is downloaded using a Web browser by going to the IBM Storage Area Network Data Gateway Web site:

<http://www.storage.ibm.com/hardsoft/products/sangateway/support/form1.htm>

This will take you to a registration window. Enter the required information and select **Submit Information**. A license agreement window is shown, and once reviewed, select **I agree**. The Download Main Page window will load. Then select the specific operating system platform. Review the readme.txt file for the latest information and instructions, before installing.

This Web site also contains all the latest firmware for the SAN Data Gateway and supported host bus adapters.

The StorWatch SAN Data Gateway Specialist software file is a self-extracting file. Once it has been downloaded, execute or run the file and it will automatically load onto your computer.

Starting the Specialist

To start the Specialist, the server must be started first, and then the client can be launched. Figure 4-4 provides an example of the StorWatch SAN Data Gateway Specialist with server and client loaded onto the same Windows NT computer.


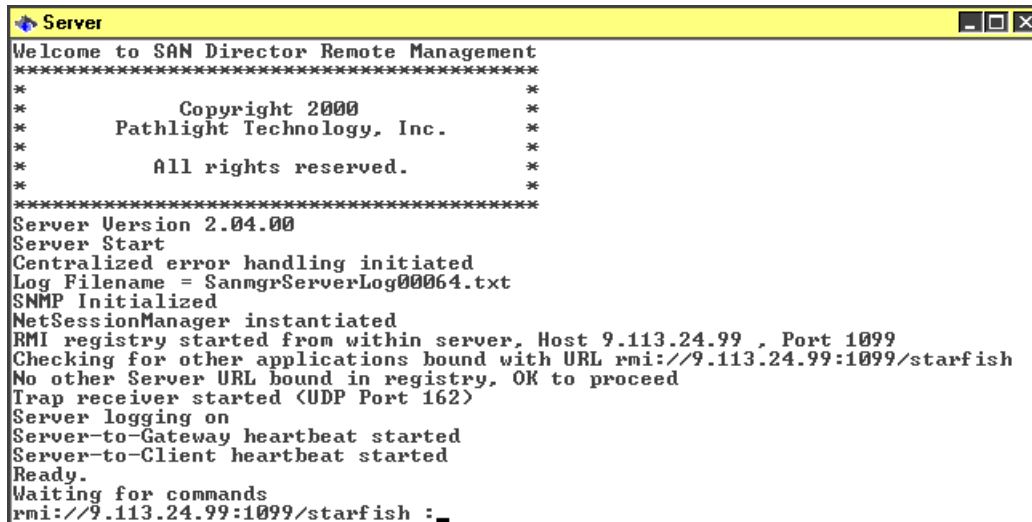


Figure 4-4 StorWatch SAN Data Gateway Specialist startup

Once the server has been launched, you should see a window similar to Figure 4-5.

A screenshot of a Windows-style window titled "Server". The window contains a text-based startup sequence for "SAN Director Remote Management". The text includes copyright information for Pathlight Technology, Inc. (2000), version 2.04.00, and a list of initialization steps: Centralized error handling, Log File, SNMP, NetSessionManager, RMI registry, and various heartbeat services. The window ends with "Ready." and "Waiting for commands", followed by a command prompt showing "rmi://9.113.24.99:1099/starfish :_".

```
Server
Welcome to SAN Director Remote Management
*****
*                               *
*      Copyright 2000           *
*      Pathlight Technology, Inc. *
*                               *
*      All rights reserved.     *
*                               *
*****
Server Version 2.04.00
Server Start
Centralized error handling initiated
Log Filename = SanmgrServerLog00064.txt
SNMP Initialized
NetSessionManager instantiated
RMI registry started from within server, Host 9.113.24.99 , Port 1099
Checking for other applications bound with URL rmi://9.113.24.99:1099/starfish
No other Server URL bound in registry, OK to proceed
Trap receiver started (UDP Port 162)
Server logging on
Server-to-Gateway heartbeat started
Server-to-Client heartbeat started
Ready.
Waiting for commands
rmi://9.113.24.99:1099/starfish :_
```

Figure 4-5 StorWatch SAN Data Gateway Specialist server

The client software can now be launched. If the server and client are not on the same PC, then a dialog box will appear to allow you to enter in the IP address of the computer that has the server software loaded. If the server and client are on the same computer you will be automatically connected to this server. After connection to the server is complete, a dialog box will appear, in which you can enter in a user name and password.

The IBM Storage Area Network Data Gateway provides a default administrator:

- ▶ User Name: StorWatch
- ▶ Password: StorWatch

The fields are case sensitive, so they must be entered in as shown above.

A new administrator account should be set up by selecting **Admin -> Add User** from the toolbar. After a new administrator account is created, then the default user StorWatch is deactivated.

Note: If a new administrator account has been created, and the password is lost, and no other account has administrator access, then you will have to contact a service representative.

4.3.2 Using the StorWatch SAN Data Gateway Specialist

Once you are logged in to the Specialist, you must now connect to the SAN Data Gateway. At this point, a dialog box should appear, requesting the IP address of the SAN Data Gateway. As it connects, it will download the information from the SAN Data Gateway and be presented on your window.

If a dialog box does *not* appear automatically, select **Tools-> Connect SAN Data Gateway or SAN Data Gateway Router** from the toolbar. This can also be used to connect to several Gateways or Routers from a single client session.

In Figure 4-6, we show the initial view once a connection to a SAN Data Gateway is established and the data has been downloaded.

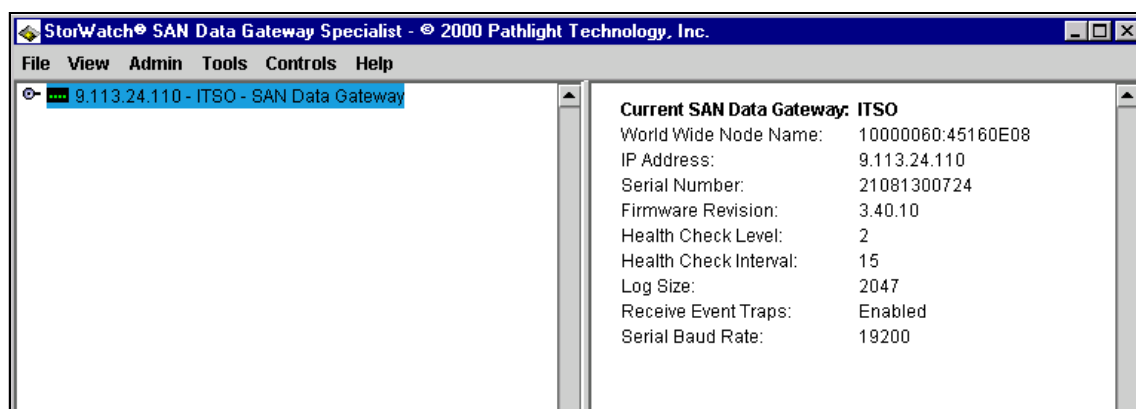


Figure 4-6 StorWatch SAN Data Gateway Specialist initial view

The left side of the window shows the SAN Data Gateway unit we are connected to, and the right side provides product data information. You will also notice that the toolbar will have options available that were previously greyed out. You can now connect to another SAN Data Gateway, disconnect from a SAN Data Gateway, enable and access the Zoning and VPS features, restart the Gateway, and also refresh the data to your window by downloading it again.

These options become available when a SAN Data Gateway is highlighted. As you begin to add SAN Data Gateway systems or drill-down into a particular Gateway by selecting and highlighting different channels or ports, different options will become available and other options will become greyed out and unavailable. Be aware of what system, channel, or port is highlighted as you move through the toolbar.

As we can connect to several SAN Data Gateway systems from one client session, select the particular Gateway you want and it will be highlighted in blue, as shown in Figure 4-7.

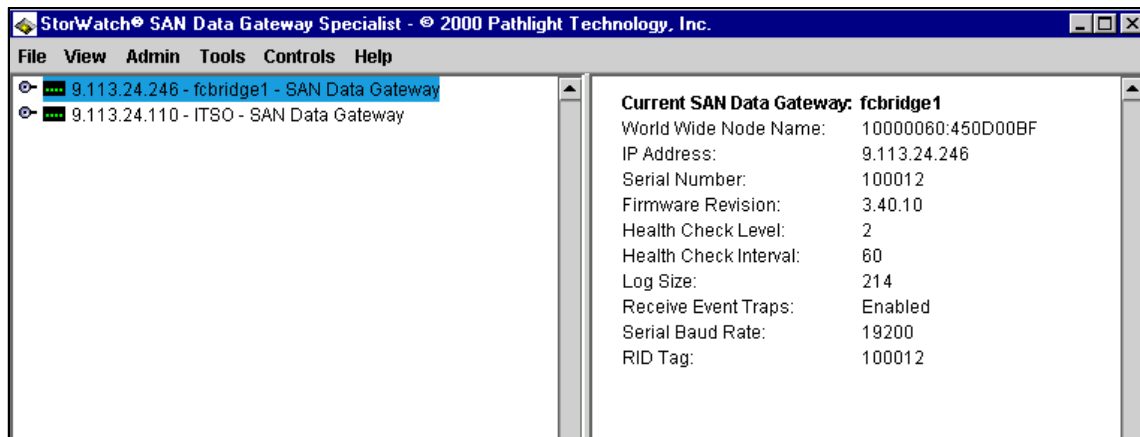


Figure 4-7 Selecting from multiple SAN Data Gateways

On the left-hand side of the highlighted Gateway, there is a small key, and by selecting this, it expands the view to show you all SCSI ports and installed Fibre Channel ports. For example, Figure 4-8 shows a Gateway with four SCSI ports and two Fibre Channel ports.

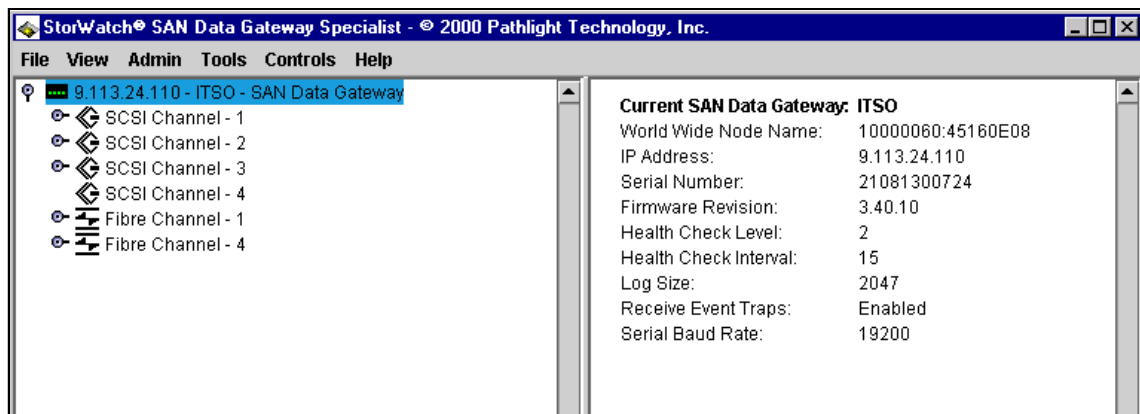


Figure 4-8 Expanded Gateway view

SCSI channels 1, 2, and 3 and Fibre Channel ports 1 and 4 also have a key on the left-hand side to depict that there are devices attached. By selecting a key, you will now expand the tree, as seen in Figure 4-9, and view the different disk devices attached.

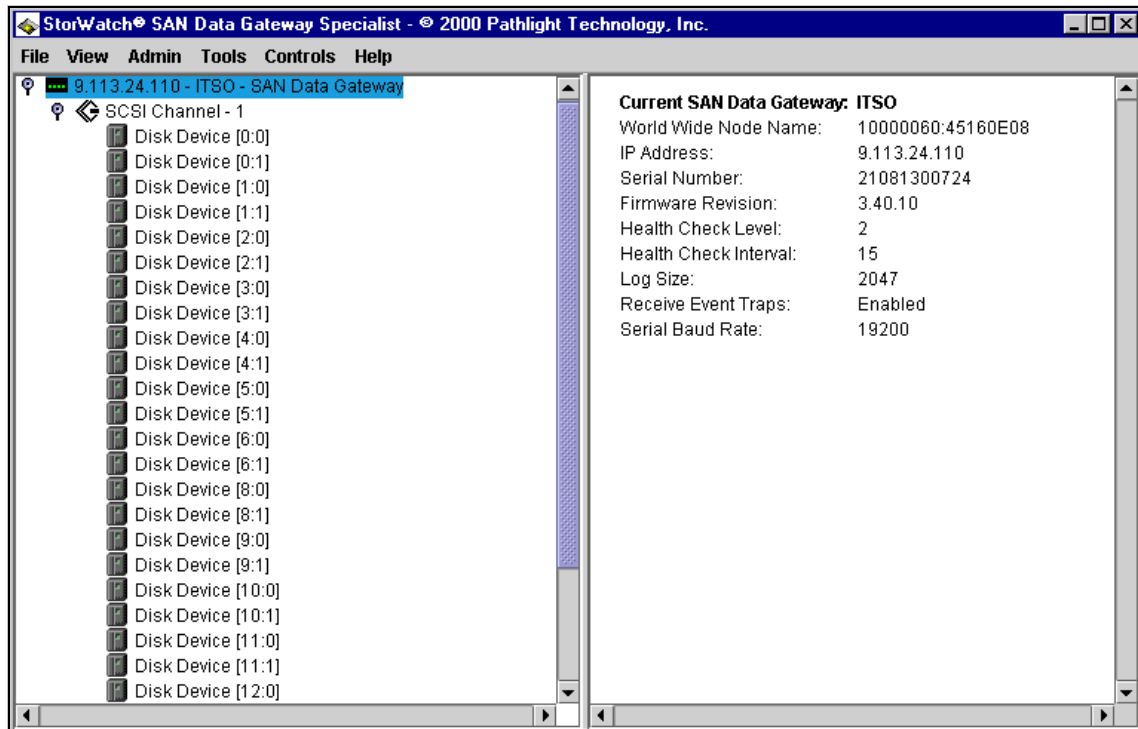


Figure 4-9 SCSI channel expanded view

You can also select and highlight each SCSI Channel. You will notice that as you do this, the information window on the right side will provide data that is unique to that SCSI channel, as shown in Figure 4-10.

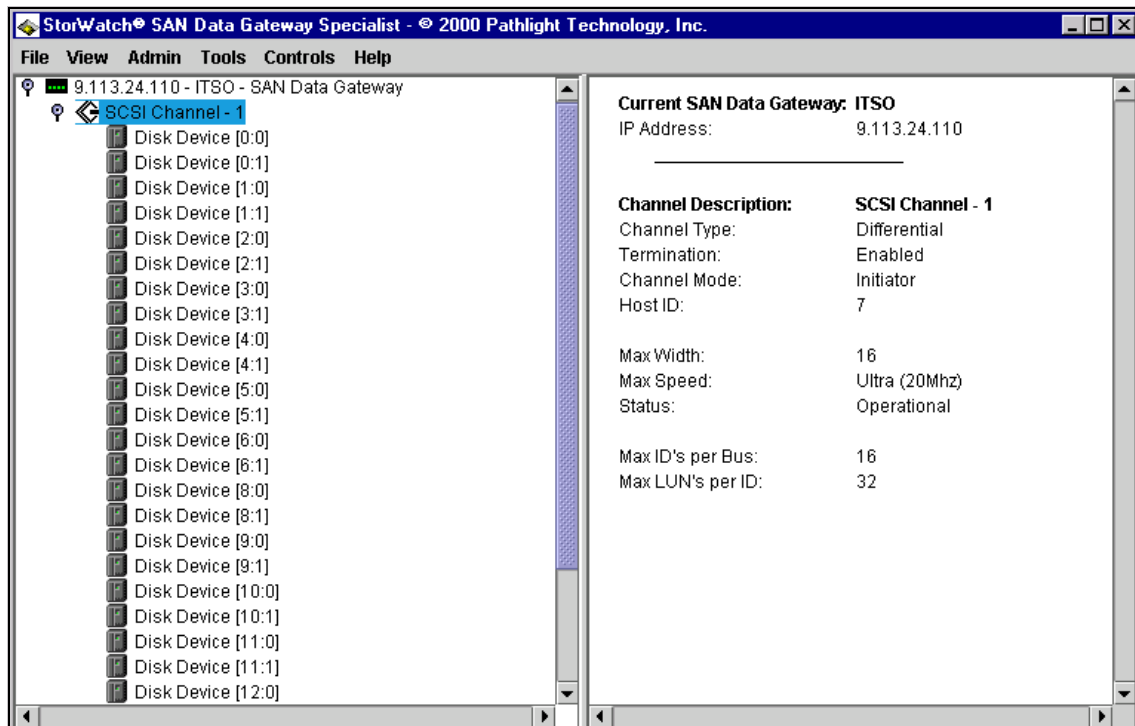


Figure 4-10 SCSI channel data

Information pertaining to a particular disk device is shown in Figure 4-11.

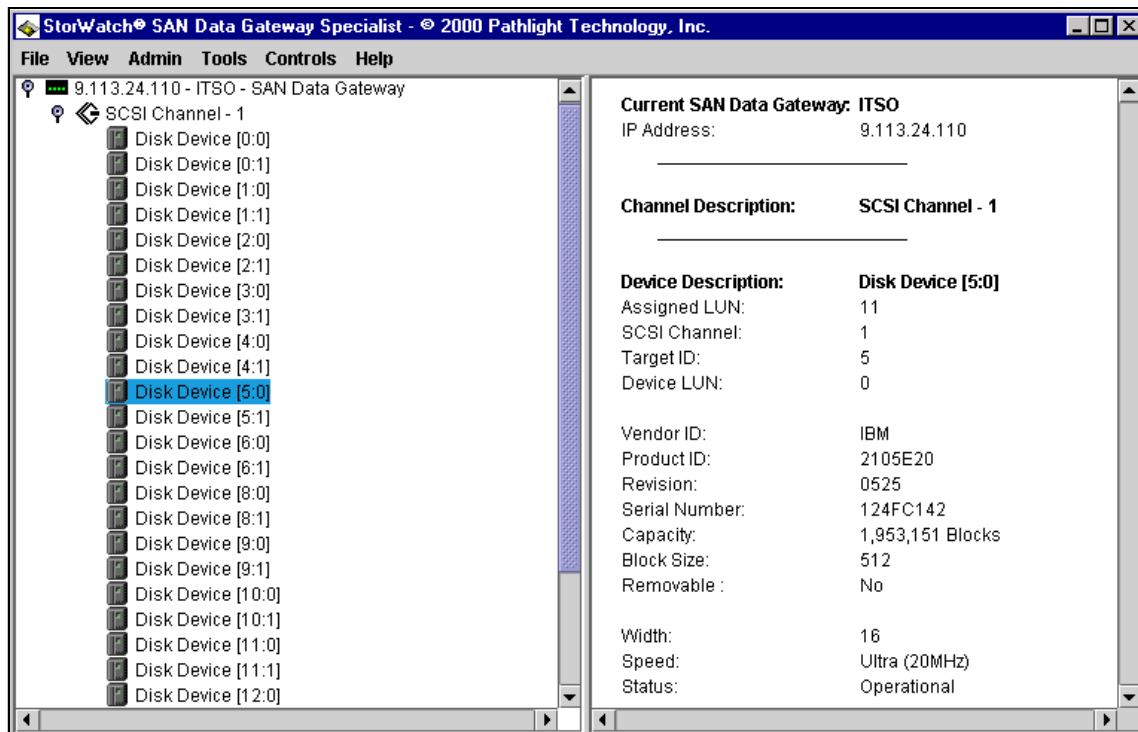


Figure 4-11 Disk device data

You can perform the same drill-down on the Fibre Channel host as we did with the SCSI channel and disk devices. Select one of the Fibre Channel port connections, as shown in Figure 4-12, and its data will be shown on the right-hand view pane.

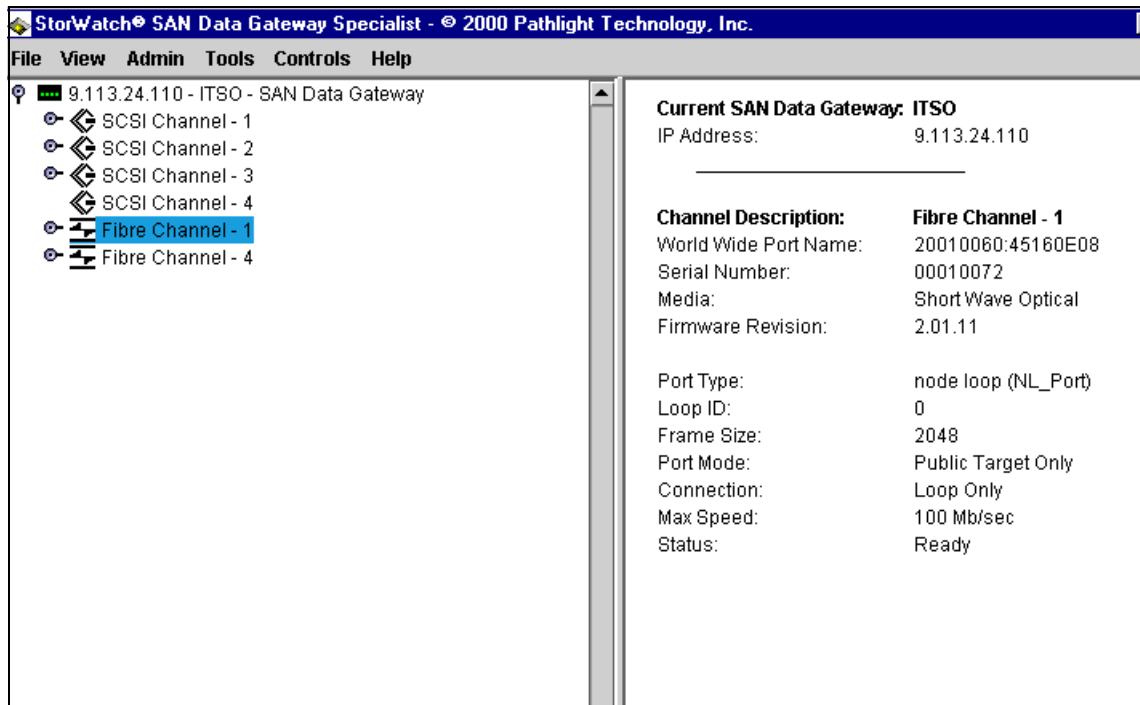


Figure 4-12 Fibre Channel port data

By selecting the key to the left, you can expand the tree and select the host system attached to that port. Figure 4-13 shows the detail on the specific host.

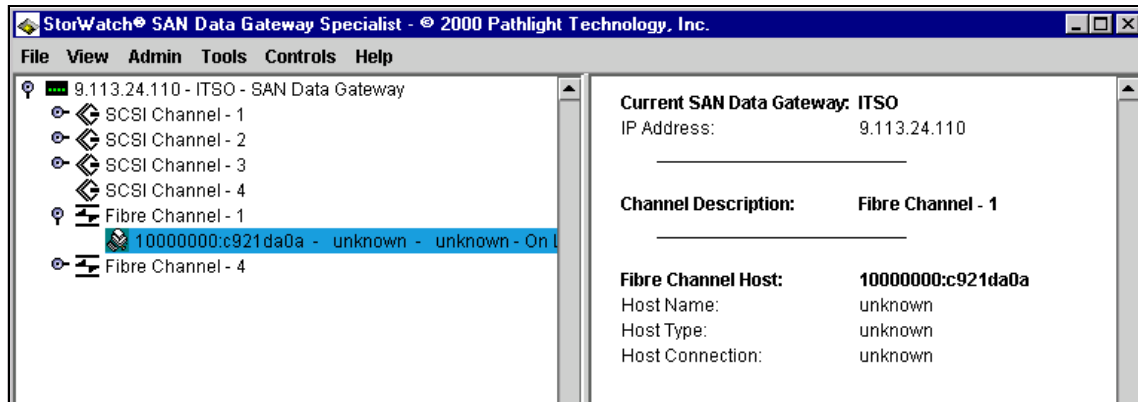


Figure 4-13 Fibre Channel host data

As you select and highlight the different ports or devices, there are different options available from the top toolbar. If an SCSI channel is highlighted, select **Controls** from the toolbar. You will notice that all options are grayed out except for **SCSI Channel**. Once selected, a dialog box will appear, as shown in Figure 4-14, and display the settings for the SCSI channel.

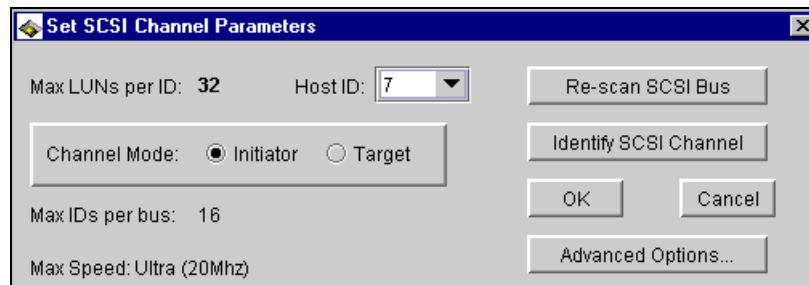


Figure 4-14 SCSI channel parameters

Selecting the **Advanced Options** button displays a dialog box, as shown in Figure 4-15. These settings are not typically changed and may disrupt normal operations, so you should consult a service representative before changing any Advanced Options.

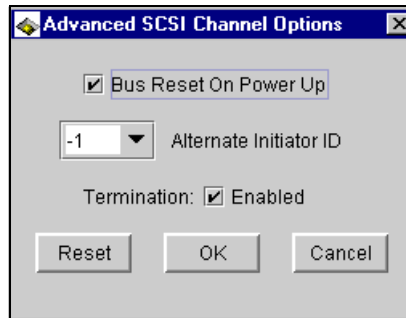


Figure 4-15 Advanced SCSI parameters

Fibre Channel parameters are displayed in a similar fashion. Highlight a Fibre Channel port and select **Controls** from the toolbar, you will notice that now all options are grayed out except the **Fibre Channel** option. By selecting this, a dialog box will display the parameters that can be changed for the Fibre Channel port selected. If any of the settings, as shown in Figure 4-16, are changed, the SAN Data Gateway must be restarted.

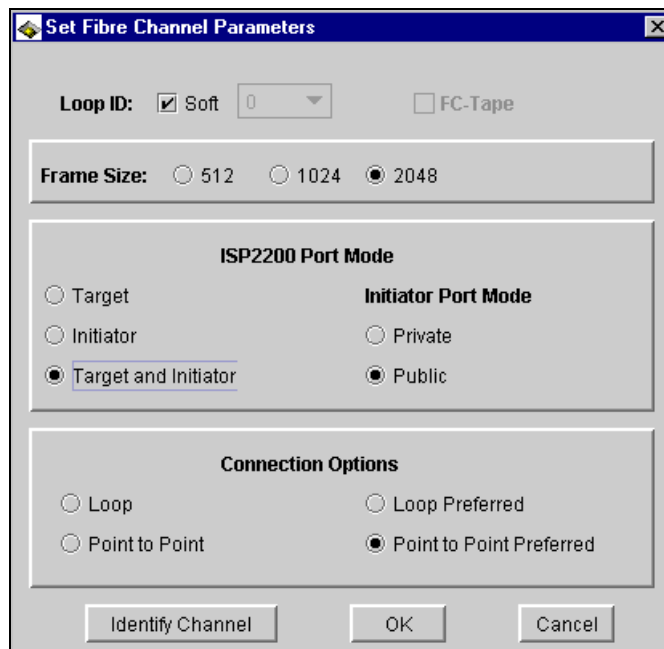


Figure 4-16 Fibre Channel port parameters

A detailed description of the SCSI and Fibre Channel settings can be found in the *IBM Storage Area Network Data Gateway Installation and User's Guide*, SC26-7304.

4.3.3 Upgrading the firmware

New versions of the SAN Data Gateway firmware can be downloaded from <http://www.storage.ibm.com/hardsoft/products/sangateway/gatewayspec.htm>

This site has a link to a Downloads section from where the latest versions of the SAN Data Gateway can be downloaded. You will first have to complete a registration form and then agree to the IBM Terms and Conditions.

Once you have completed this section, you will find a screen similar to that shown below in Figure 4-17.

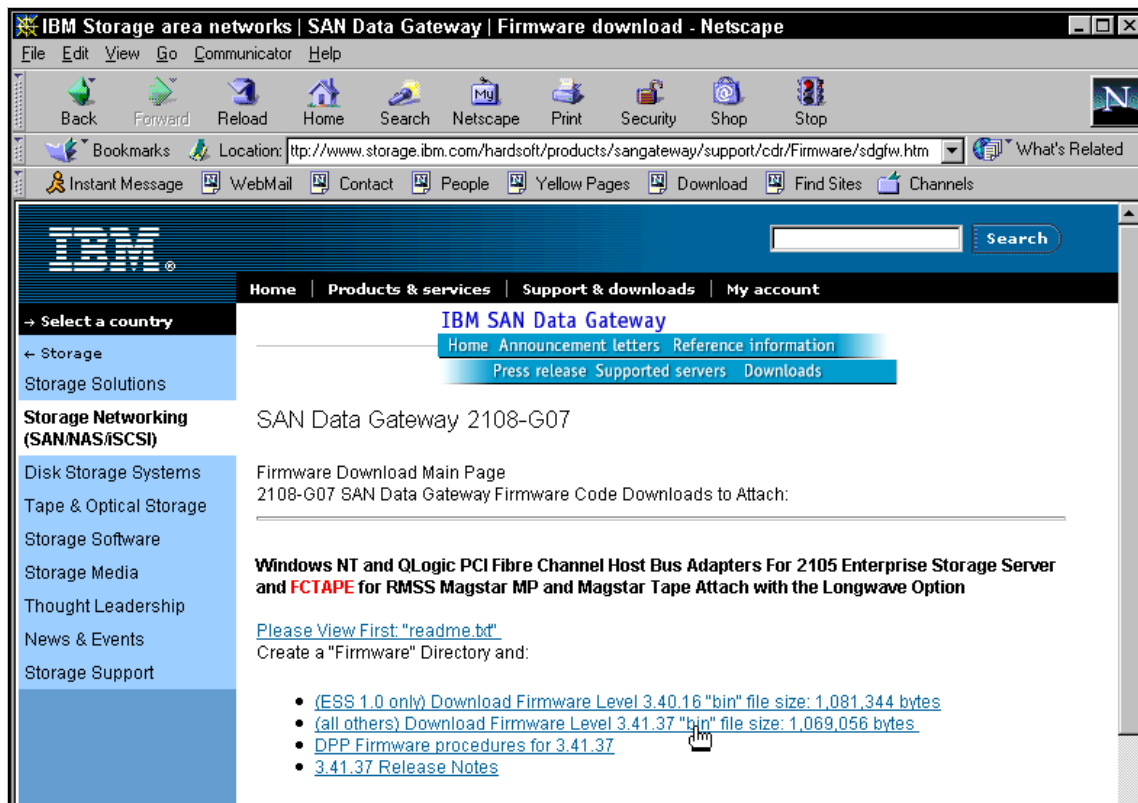


Figure 4-17 Downloading the SAN Data Gateway firmware

From this screen, we selected the Windows NT version, and we will save this into a directory for downloading into the SAN Data Gateway at a later stage.

From the main screen of the SAN Data Gateway Specialist, you can check the current Firmware Revision level, as shown in Figure 4-18.

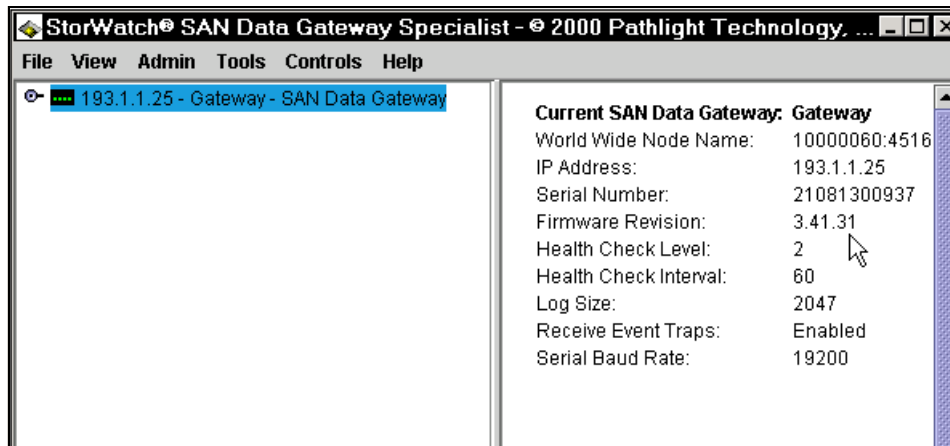


Figure 4-18 SAN Data Gateway Firmware Revision Level

New versions of the SAN Data Gateway can be downloaded, as shown in Figure 4-19, by selecting the **Controls** option and then selecting the **Update Firmware** option.

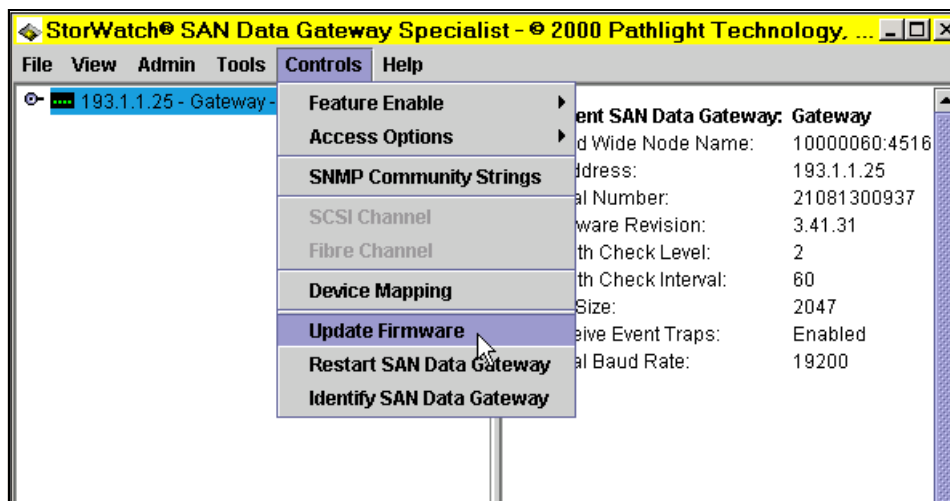


Figure 4-19 Updating the SAN Data Gateway firmware

The SAN Data Gateway Specialist will then prompt for the location of the new firmware as shown in Figure 4-20. This is the file that was downloaded from the SAN Data Gateway Web site previously.

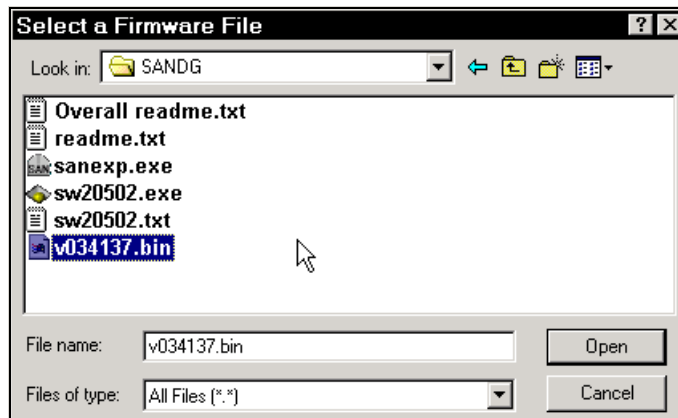


Figure 4-20 Specifying location of the firmware

Downloading the firmware into the SAN Data Gateway is a disruptive process, so the Specialist displays a warning message as shown in Figure 4-21.




Figure 4-21 Warning message prior to downloading the firmware

Once all I/O activity is stopped, you can click **Yes** to continue the process.

Once the firmware process has been started, the SAN Data Gateway Specialist displays a message as shown in Figure 4-22.

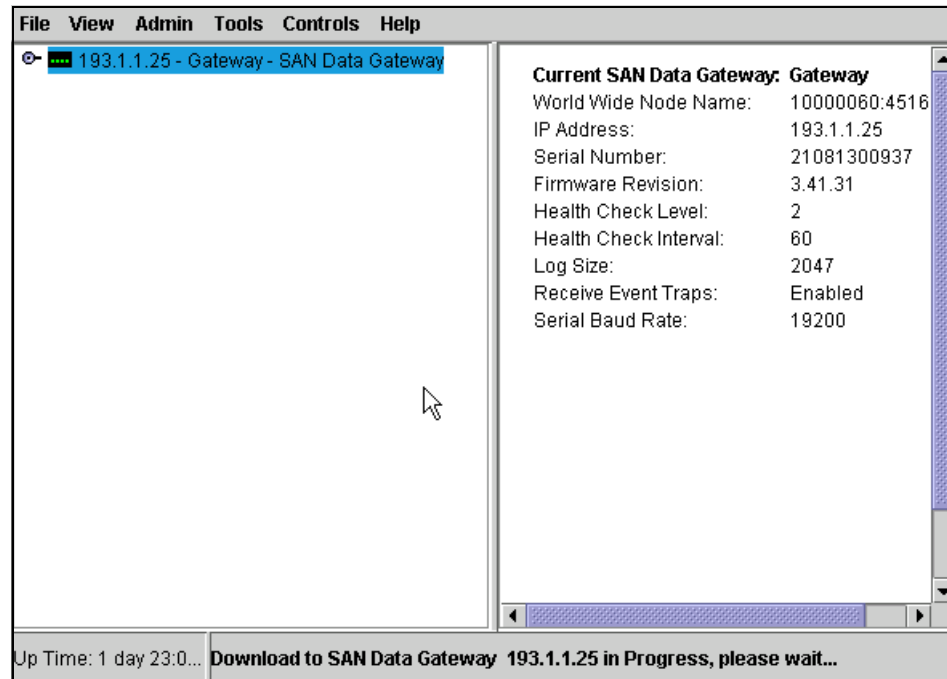


Figure 4-22 Download in progress

This indicates that the firmware is being downloaded to the SAN Data Gateway.

The Specialist will then give the option to restart the SAN Data Gateway as shown in Figure 4-23.




Figure 4-23 Message prior to restarting the SAN Data Gateway

As restarting the SAN Data Gateway is a disruptive process, the Specialist issues a warning as shown below in Figure 4-24.




Figure 4-24 Warning message prior to restarting the SAN Data Gateway

The Specialist will now return to the main screen, and a message is displayed at the bottom of the screen as shown in Figure 4-25 indicating that the Gateway is in the process of restarting.

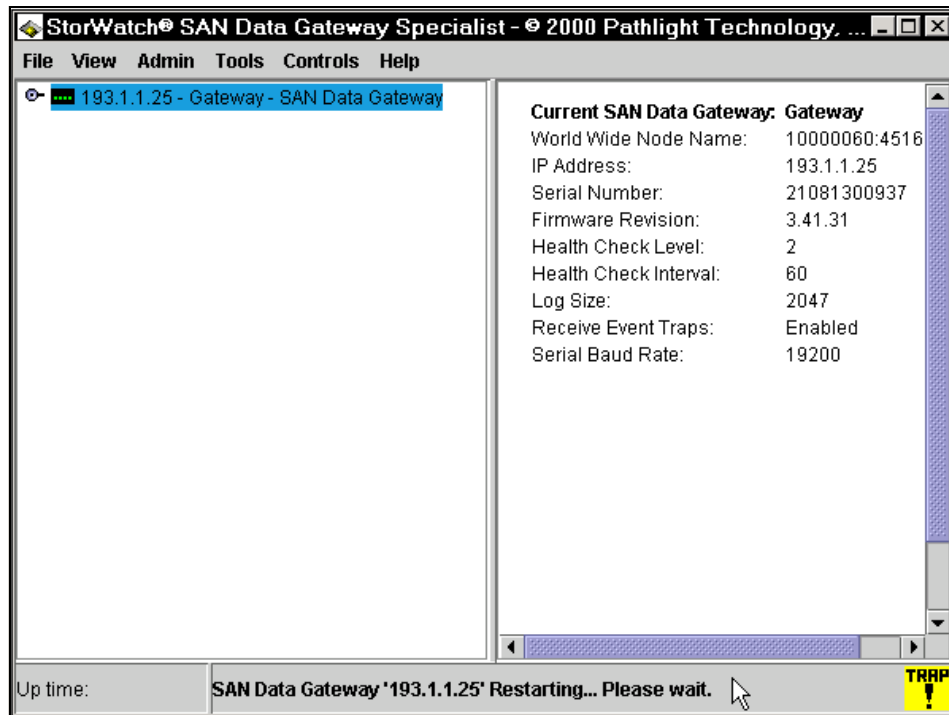


Figure 4-25 SAN Data Gateway now restarting

Once the restart process is completed, the Specialist displays the prompt shown in Figure 4-26 before refreshing the display.




Figure 4-26 Restart completed message

Once the Specialist has finished refreshing, we can see that the firmware revision level has been updated successfully, as shown in Figure 4-27.

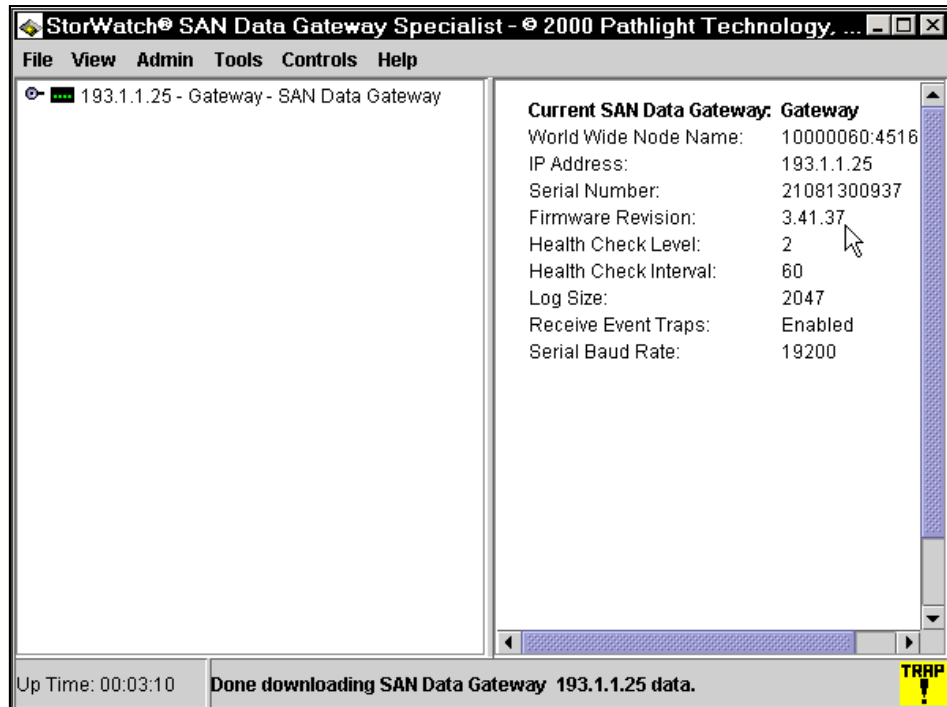


Figure 4-27 New firmware revision level

4.4 SCSI devices

The four SCSI ports on the SAN Data Gateway support Differential Ultra Wide SCSI devices. It will automatically negotiate speed for wide or narrow bus width devices as well as standard, fast and, ultra speeds. The SAN Data Gateway provides a termination on each of its SCSI buses. The termination can be separately disabled if so desired from the StorWatch SAN Data Gateway Specialist window.

4.4.1 LUN support

The SAN Data Gateway provides support for up to 256 LUNs. Each SCSI channel supports up to 15 SCSI targets and up to 32 LUN per target. This is subject to the overall total of 256 available. The first LUN (LUN 0) is used for the Gateway for control and command purposes. That leaves 255 allowable LUN addresses to be used for attaching SCSI devices to the four SCSI ports.

Note: Do not attach more than 255 LUNs to the SAN Data Gateway. Exceeding the 255 limit will cause unpredictable results.

The SCSI devices must be previously set up in your host prior to attaching to the SAN Data Gateway. For example, if an IBM Enterprise Storage Server is to be used as the target device, then all the volume or LUN assignments must be completed. Also, the volumes should be assigned to the appropriate SCSI port in the ESS. Attach the SCSI cable from the SCSI device to a SCSI port on the Gateway.

4.4.2 Device discovery

Once attached, the SAN Data Gateway can be restarted or powered on to perform a discovery of the attached devices. If this was done after the **initializeBox** command from the service terminal, then the initial discovery will be carried out in a specific order. The buses are scanned in order from 1 to 4 and each bus is scanned from SCSI target 0 to 15 and LUN 0 to 32 for each ID. As a device is discovered it is assigned a specific LUN number by the SAN Data Gateway. This numbering will begin at LUN number 1, zero is reserved for the SAN Data Gateway control, and continues sequentially as devices are discovered.

The SAN Data Gateway must add this extra layer of addressing as the host is no longer directly attached to the disk devices but will see a single target ID, the SAN Data Gateway. In a regular SCSI environment with the host attached directly to the SCSI device, the host is able to recognize the target and LUN ID of that device. Since we add the SAN Data Gateway in between the host and the device, the host is no longer directly attached to the SCSI device to recognize its target or LUN ID.

4.4.3 Assigning LUN IDs

The Gateway assigns LUN IDs to all the attached devices, up to the maximum of 255, and it creates a map of the actual disk device target and LUN ID to its own assigned LUN ID. This map is stored in nonvolatile memory within the Gateway.

4.4.4 Adding SCSI devices

New SCSI devices can be added at any time. If a new device is added, for example, another volume in the IBM Enterprise Storage Server has been assigned to a SCSI port on a San Data Gateway, the SCSI buses must be re-scanned to detect the new device. This can be done from the service terminal using the **scsiRescan** command or from the StorWatch SAN Data Gateway Specialist. If using the Specialist, select and highlight the SCSI bus that the device has been assigned to, and select **Controls-> SCSI Channel-> Re-Scan SCSI Bus**. As it rescans, the SAN Data Gateway will assign the next available LUN address to the new device. Refresh the data on your Specialist window by selecting **View-> Refresh SAN Data Gateway**.

Once a device has been assigned a LUN ID by the Gateway, it will maintain that ID since it was written into the device map. This is useful in case the device is lost or needs to be replaced. Remove the old device, set the new device to the same SCSI bus target and LUN as the old device and attach it to the same channel. You must rescan the SAN Data Gateway SCSI bus for it to update its data. You will notice that the replaced device has kept the same assigned LUN.

If a device is no longer attached or no longer required, then the assigned LUN is also no longer required. To free up this assigned LUN ID, you can issue the **mapWinnowDatabase** command from the service terminal. The existing devices will maintain their assigned LUN IDs.

If a complete rebuild of the SCSI device map is desired or required, this can be done from the service terminal only. You would issue the **mapRebuildDatabase** command. This command deletes the existing device map and the SAN Data Gateway will restart. When it is rebooted, a new scan of the SCSI buses is done as if the system was brand new, and LUN IDs will be assigned as described earlier.

When this command is issued, you will also be given an option to clear the Virtual Private SAN (VPS) access settings. Because this allows host access to specific LUNs and by issuing this command, we may change the assigned LUN ID; it is recommended that you always say **Yes** to this option. If not, a host may access a volume that you do not want it to access and be restricted from a volume that it had access to previously. The system administrator must rebuild the VPS settings, if enabled, to allow host access to the desired LUNs.

4.5 SAN Data Gateway access options

The SAN Data Gateway includes two features that are helpful in providing control and security of host access to SCSI devices:

- ▶ Zoning
- ▶ Virtual Private SAN (or LUN masking)

4.5.1 Zoning

Channel zoning is a feature included with the IBM Storage Area Network Data Gateway. Zoning allows you to restrict access between SAN Fibre Channel connections and SCSI channels. The default settings allow all SAN connections to access all SCSI channels.

Configuring zoning

To configure zoning, and you must be an administrator, select **Control -> Access Options -> Channel Zoning** from the toolbar. Figure 4-28 shows a zoned SAN Data Gateway.

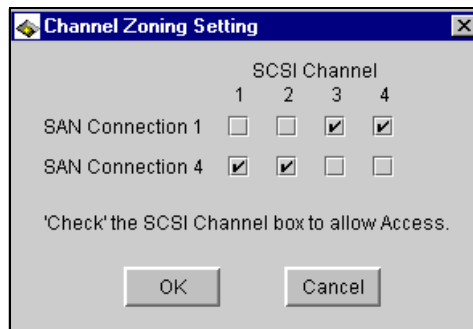


Figure 4-28 IBM Storage Area Network Data Gateway channel zoning

A check mark will allow access, and as shown in Figure 4-28, you can see that SAN connection 1 has access to SCSI channels 3 and 4, but not to SCSI channels 1 and 2. SAN connection 4 has access to SCSI channels 1 and 2, but not to SCSI channels 3 and 4. To change the settings, click any box and the check mark will toggle on and off. All combinations are possible. Once the desired settings are selected, click **OK**. For the new zone settings to take effect, the SAN Data Gateway must be restarted.

4.5.2 Virtual Private SAN

The IBM Storage Area Network Data Gateway also provides LUN masking through a feature called Virtual Private SAN (VPS). This provides a granular restriction of host access to specific LUNs while utilizing and sharing the same connectivity paths through the SAN Data Gateway.

VPS keeps track of Fibre Channel hosts by using their unique World Wide Name. In this way, if a switch is attached, the Gateway will also keep track of it, because the switch also has a unique World Wide Name. VPS also includes a host registration service so that when a host is attached, its unique WWN is registered to the Gateway to provide information on the host and its status.

Enabling VPS

The Virtual Private SAN feature has to be enabled, because it comes disabled from the factory. To enable it, make sure the desired SAN Data Gateway is highlighted, if more than one appears in your Specialist window. Select **Controls -> Feature Enable -> Virtual Private SAN**, as shown in Figure 4-29.

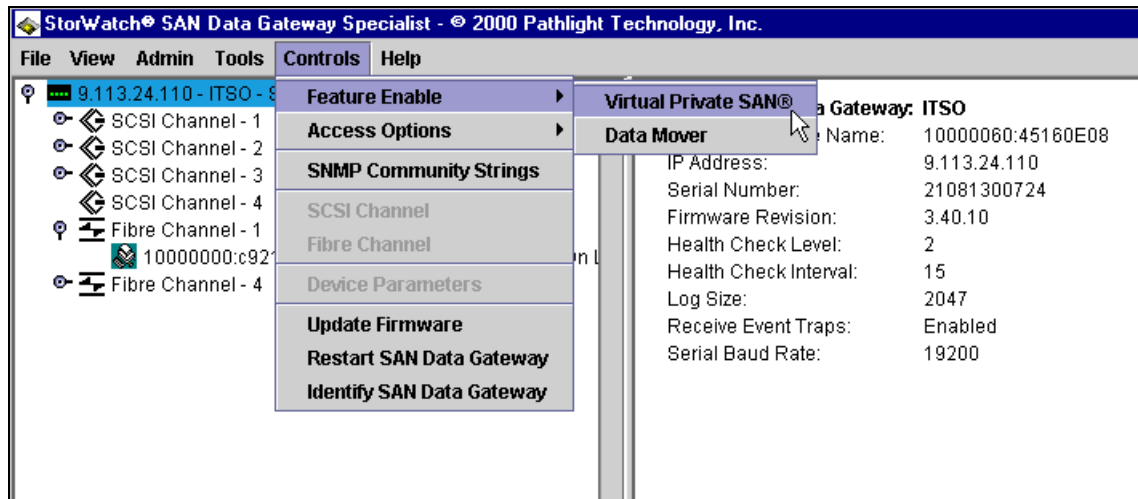


Figure 4-29 Enabling Virtual Private SAN

A dialog box appears requesting the entry of a license key number. For units with serial numbers lower than 1300600 you are required to get a license key that is matched to the serial number of the SAN Data Gateway. For units with serial numbers higher than 1300600, the VPS feature is bundled into the product. Type the word **enable** and the feature will now be enabled.

Host registration

With VPS enabled, the host registration service is also available. This service is provided to simplify VPS configurations. It provides the host name, host type, host connection, and the unique WWN of the Fibre Channel host bus adapter. Host information is sent over the Fibre Channel connection to the Gateway.

To obtain this information, a program is loaded on each host. This program is found, and can be downloaded at the same Web site that the StorWatch SAN Data Gateway Specialist was downloaded from.

<http://www.storage.ibm.com/hardsoft/products/sangateway/support/form1.htm>

After completing the registration and license agreement, the Download Main page is displayed. Select the operating system software subheading and look for **HOSTSW**. View the read.me file for the latest information. Download the software and install it onto the host that will be attaching to the SAN Data Gateway. Follow the instructions provided in the read.me file.

For Windows NT, the file is a self-executing file, so it can be executed or run as any *.exe file. As it runs through the install, make sure to select **SAN Explorer** and **VPS Registration Service**, as shown in Figure 4-30.




Figure 4-30 Loading VPS Registration software on Windows NT

Once installed, the service runs automatically and does not require further administrator interaction.

Communicating with the Gateway

Once the host is attached to the Gateway and restarted, the registration service will communicate to the Gateway. The data shown on the Gateway will have to be refreshed by selecting **View-> Refresh SAN Data Gateway**. This will cause the updated data to be shown in the Specialist window.

The registration service will re-register the host to the SAN Data Gateway at a default of 15 minute intervals. This interval can be changed if so desired.

Previous to enabling the VPS feature, you will have seen that the Specialist displayed a window similar to Figure 4-31. There is no key beside the Fibre Channel ports, indicating that no host is recognized.

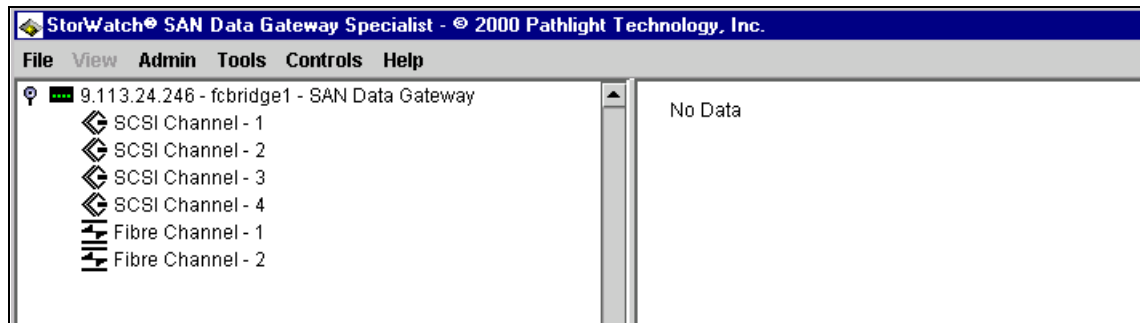


Figure 4-31 Specialist display without VPS enabled

After the VPS is enabled and a host has registered with the Gateway, all its information will load automatically into the database of the Gateway. Figure 4-32 shows, in the right-hand side view pane, the host name, host type, HBA and connection information, and the unique WWN that was obtained automatically by host registration.



Figure 4-32 Specialist after VPS enabled and host registration

If no registration software is loaded on the host, or is not available for a specific operating system, for example a Fibre Channel switch, only the WWN of the attached system will register to the VPS database, all other fields will have *unknown*. This is shown in Figure 4-33.

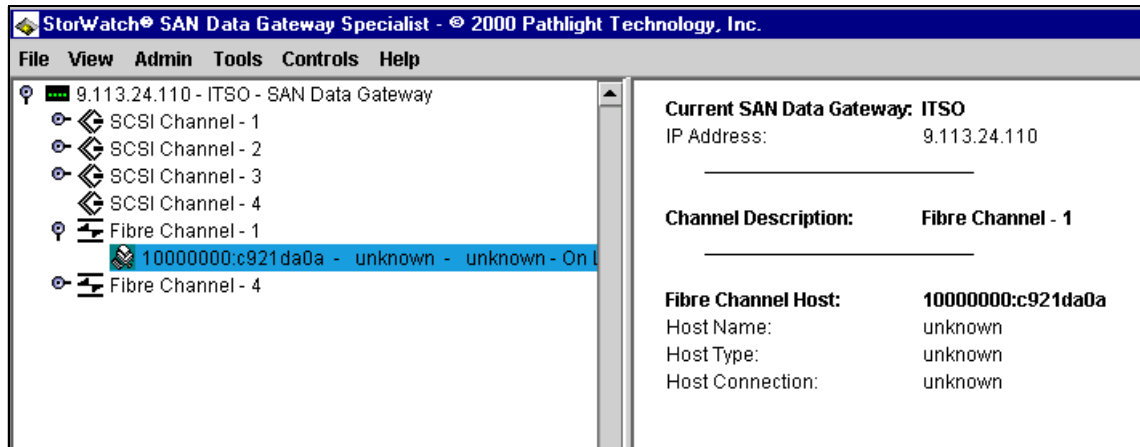


Figure 4-33 Host system with no host registration software

A diagram to depict the setup that was described, is shown in Figure 4-34. Note that as the first dual Fibre Channel port on the Gateway is used, it is numbered as input 1 and 4. A description on the Fibre Channel port numbering was shown in Figure 4-2 on page 273.




Figure 4-34 SAN Data Gateway with two hosts

Adding host and connection information

The host name and host type and connection information can be added manually by selecting the VPS feature and modifying the data. To do this, select the desired SAN Data Gateway so that it is highlighted. Select **Controls -> Access Options -> Virtual Private SAN**, and you will enter into the VPS settings window. The hosts and their information will appear on the left-hand side.

Select the host to be modified, and the information is transferred to the bottom part of the window where the modifications can take place, as shown in Figure 4-35.

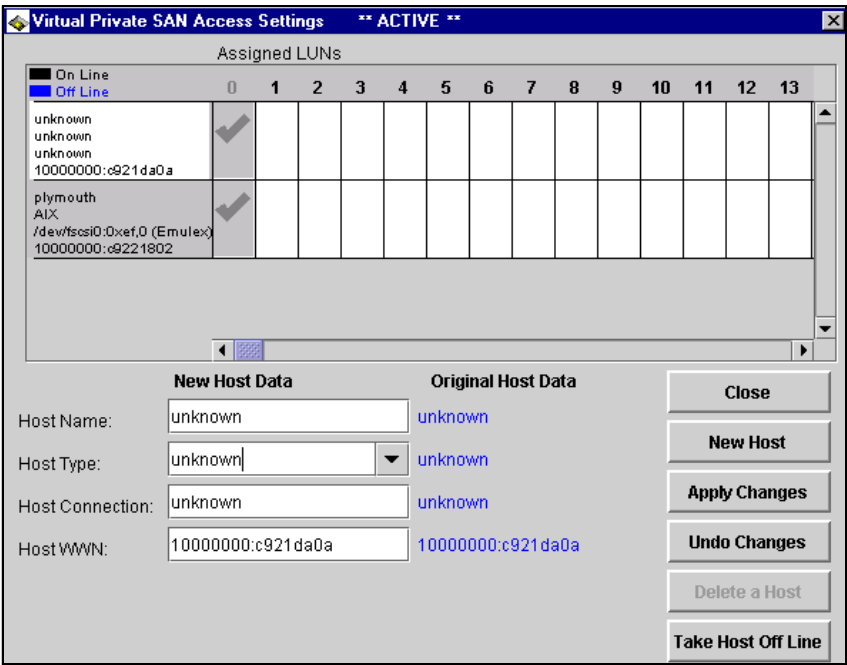


Figure 4-35 VPS host settings

The Host information can now be entered in the left-hand column labeled New Host Data. The Original Host Data column displays the values before any changes are made. Once completed, select **Apply Changes** and then **Close**.

This window can also be used to enter in host information for a host that has not yet connected to the SAN Data Gateway. Select **New Host** and then enter in the WWN and other parameters that are available. You *must*, at minimum, enter in a WWN. Also, when entering in the number, a colon “:” *must* be used to separate the first four bytes from the last four bytes.

Also, you will notice that the host description in the top can change color. If the information is in *blue*, then the host is offline. If the host information is in *black*, then it is online.

The VPS feature allows an administrator to quickly view the host information and status at a central location.

Setting up a Virtual Private SAN

Remember that the SAN Data Gateway assigns its own LUN numbers, and it does so in a sequential order. An attached disk device may have a SCSI target of 1 with a LUN ID of 4, but when recognized by the SAN Data Gateway, its LUN number assigned will be something completely different. An example is shown in Figure 4-36.

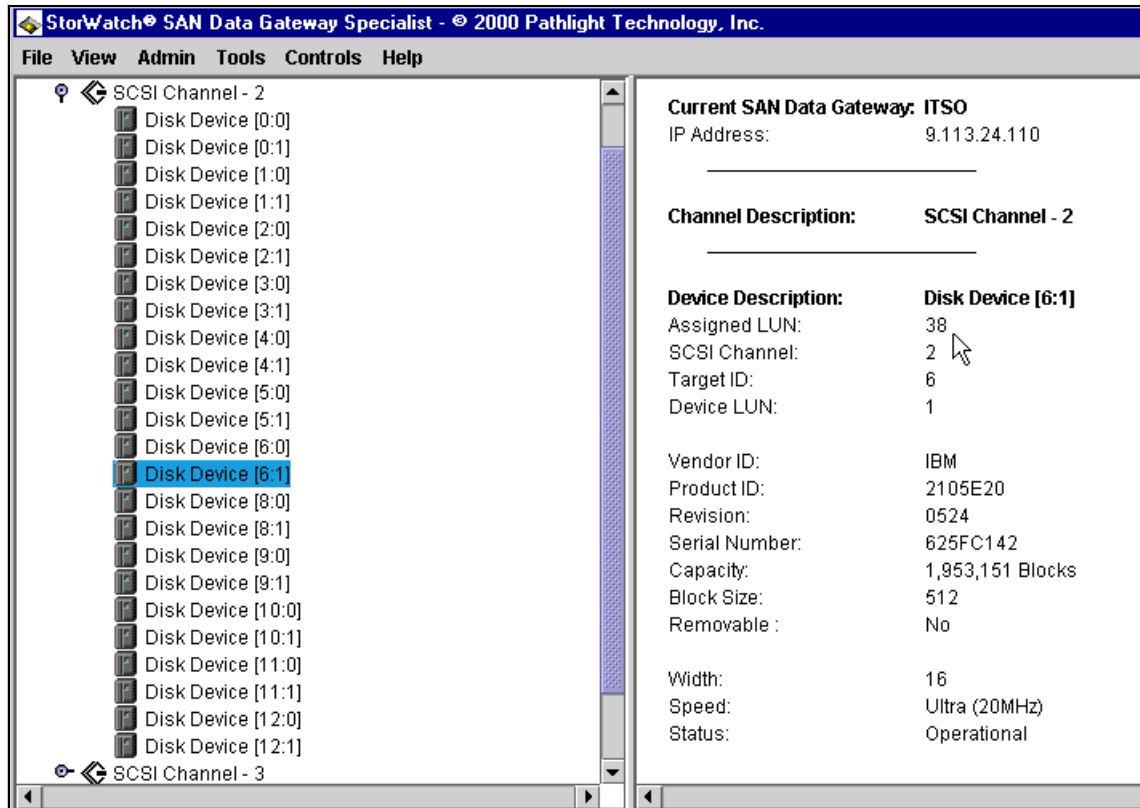


Figure 4-36 SCSI LUN assignment

In this example, the disk device attached has a SCSI target of 6 and LUN ID of 1. When it was recognized by the SAN Data Gateway, the Gateway assigned LUN number 38 to this device. This Gateway's assigned LUN number corresponds to the assigned LUN number that appears on the top of the VPS settings window, as shown Figure 4-38 on page 305.

SCSI to LUN map

Prior to setting the LUN masking, it makes sense to prepare a list that has each LUN and shows which host is allocated to that LUN. To assist in this process the Gateway has a map that provides a cross reference of the actual SCSI target and ID to the Gateway's assigned LUN number. To access this map you must access the service terminal from the serial port as described earlier in 4.2, "Installation" on page 273. Once connected, type in the command **mapShowDatabase**. The database lists not only devices that are presently connected, but also devices that have previously been connected. If a previously attached device is later reattached, it is assigned back to its previous address. Figure 4-37 shows the output returned to the service terminal.

ITS0 > mapShowDatabase						
devId	Type	Chan	tId	tLun	UID	
000	SMA	127	127	127	00000000:45160e08	
001	SCSI	001	000	000	20100000:45160e08	
002	SCSI	001	000	001	20200000:45160e08	
003	SCSI	001	001	000	20300000:45160e08	
004	SCSI	001	001	001	20400000:45160e08	
005	SCSI	001	002	000	20500000:45160e08	
006	SCSI	001	002	001	20600000:45160e08	
007	SCSI	001	003	000	20700000:45160e08	
008	SCSI	001	003	001	20800000:45160e08	
009	SCSI	001	004	000	20900000:45160e08	
010	SCSI	001	004	001	20a00000:45160e08	
011	SCSI	001	005	000	20b00000:45160e08	
012	SCSI	001	005	001	20c00000:45160e08	
013	SCSI	001	006	000	20d00000:45160e08	
014	SCSI	001	006	001	20e00000:45160e08	
015	SCSI	001	008	000	20f00000:45160e08	
016	SCSI	001	008	001	21000000:45160e08	
017	SCSI	001	009	000	21100000:45160e08	
018	SCSI	001	009	001	21200000:45160e08	
019	SCSI	001	010	000	21300000:45160e08	
020	SCSI	001	010	001	21400000:45160e08	
021	SCSI	001	011	000	21500000:45160e08	
022	SCSI	001	011	001	21600000:45160e08	
023	SCSI	001	012	000	21700000:45160e08	
024	SCSI	001	012	001	21800000:45160e08	
025	SCSI	002	000	000	21900000:45160e08	
026	SCSI	002	000	001	21a00000:45160e08	
027	SCSI	002	001	000	21b00000:45160e08	
028	SCSI	002	001	001	21c00000:45160e08	
029	SCSI	002	002	000	21d00000:45160e08	
030	SCSI	002	002	001	21e00000:45160e08	

Figure 4-37 Service terminal display of device map

The numbers on the left are the assigned LUN numbers, and note that number '0' has been assigned to the Gateway. The other columns contain the device type, the SCSI channel it is connected to on the SAN Data Gateway, and the actual target and LUN ID.

Setting host access to LUNs

To view and set host access to particular LUNs, access the Virtual Private SAN Access Settings by selecting **Controls -> Access Options -> Virtual Private SAN**. This window will show all the hosts that have registered to the Gateway.

To allow a host to access a particular LUN, place a check mark in the row that corresponds to the host. To disable access, the square must be clear, without a check mark. The check mark is toggled on and off by clicking in each square. Figure 4-38 provides an example of LUN masking.

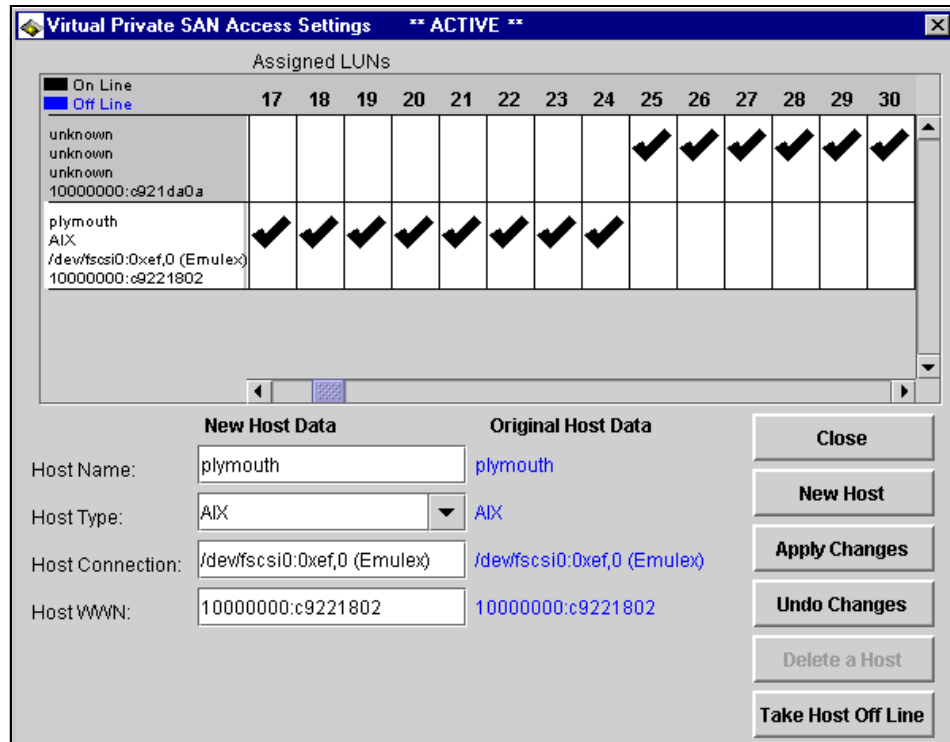


Figure 4-38 Setting LUN masking

The host *plymouth* has access to LUNs 17 to 24 but does not have access to LUNs 25 to 30. Conversely, the host at the top has access to LUNs 25 to 30, but does not have access to LUNs 17 to 24. Using the scroll bar on the bottom, we can scroll through all 255 available LUNs in the SAN Data Gateway, and enable or disable access by adding or clearing the check mark. The scroll bar on the right-hand side allows us to scroll through the different hosts. In this example, there are only two hosts, but there can be several more.

Once completed, select **Apply Changes** and then **Close**. The host system may now have to be restarted or some other method used to rescan the bus for the host to detect that it now has access to new devices.

Any combination is allowed, so if the same LUN is to be shared by two different hosts, a check mark for that LUN must be set for both. If this is the case, the host systems must have a device sharing software installed to control access to the disk device for data integrity.

As each Fibre Channel port can support up to 8 eight different hosts, there can be up to 48 hosts attached to the SAN Data Gateway. They can all share the same four SCSI channels. By using the Virtual Private SAN feature, you can ensure that only the LUNs you desire a certain host to use will be accessed and that no other host will access them.

4.5.3 Combining Zoning and Virtual Private SAN

If Virtual Private SAN is enabled and LUNs have been allocated to specific hosts, then zoning is not necessary or required. The Channel Zoning window can remain at the default settings with all Fibre Channel ports accessing all SCSI channels.

However, they can work in combination to add an extra level of control and security. If zoning is added, then VPS can only control the LUNs that are included in its access zone. Figure 4-39 shows a zoned SAN Data Gateway.

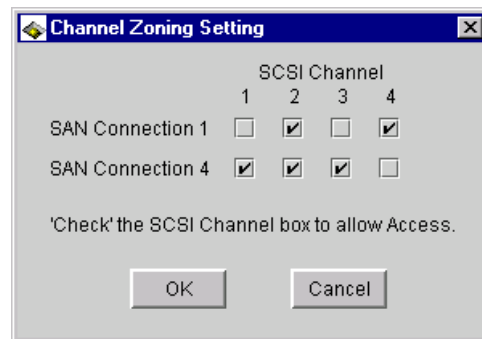


Figure 4-39 Combining channel zoning and VPS

SAN connection 1 has access to the LUNs on SCSI channels 2 and 4. VPS will control access to LUNs on SCSI channels 2 and 4 for SAN connection 1. Since SAN 1 is *not* zoned for SCSI channel 1 and 3, it will not reach any LUNs on these channels. Even, if a LUN on SCSI 3 has been enabled for access to a host on SAN 1 in the VPS settings window, the host will not see that LUN because of the zoned settings. The same is true for SAN connection 4. A host connected here will access LUNs only on SCSI channels 1, 2, and 3, but not channel 4, regardless of whether the VPS settings will allow it.

4.6 Adding Fibre Channel fabric components

The SAN Data Gateway is an integral component in a storage network. As such, you can also attach other SAN fabric components to the Gateway to increase the connectivity options in a SAN. Hubs and switches can be easily added, and allow many more systems, local or remote, to access the SCSI devices on the Gateway.

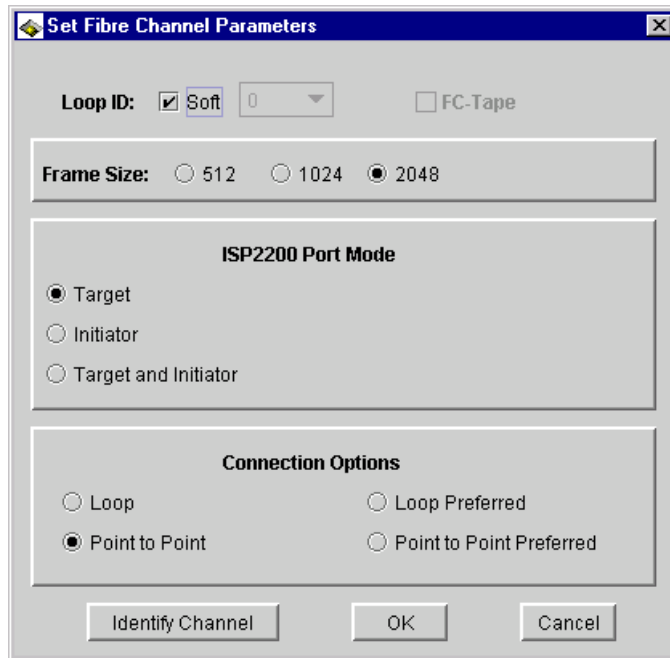
As switches provide more flexibility, and hubs are mainly used to extend distances, the following discussion will focus on switch implementation, rather than hubs.

4.6.1 Connecting an IBM SAN Fibre Channel Switch

The IBM SAN Fibre Channel Switch provides either an 8 port or 16 port switch that can also be cascaded.

Allowing fabric connection

Before connecting a switch to the SAN Data Gateway, there is a setting that must be checked to allow proper fabric connection. From the SAN Data Gateway Specialist, select the Fibre Channel port that will have the switch connected. Select **Control -> Fibre Channel**, and the Set Fibre Channel Parameters window appears, as shown in Figure 4-40.



The dialog box titled "Set Fibre Channel Parameters" contains the following settings:

- Loop ID:** A checked checkbox labeled "Soft" followed by a dropdown menu showing "0". An unchecked checkbox labeled "FC-Tape" is to the right.
- Frame Size:** Three radio buttons: "512", "1024", and "2048". The "2048" button is selected.
- ISP2200 Port Mode:** Three radio buttons: "Target", "Initiator", and "Target and Initiator". The "Target" button is selected.
- Connection Options:** Four radio buttons arranged in two columns. The left column has "Loop" and "Point to Point" (selected). The right column has "Loop Preferred" and "Point to Point Preferred".
- Buttons:** "Identify Channel", "OK", and "Cancel" at the bottom.

Figure 4-40 Fibre Channel port setting for switch attachment

By default, the setting in the **Connection Options** box will have **Loop** enabled. For switch or fabric connection, select **Point to Point**, and then select **OK**. The SAN Data Gateway must be restarted for the change to take effect.

Note: **Connection Options** box does not appear, the Fibre Channel module installed will support loop connection only and will not support fabric connection. The module must be replaced or select a port that will support fabric connection.

A Fibre Channel cable can be connected from the switch to the port on the Gateway. The data to the Specialist must be refreshed by selecting and highlighting the Gateway and then select **View -> Refresh SAN Data Gateway**.

Switch registration

With VPS enabled, the switch will register with the database the WWPN of the port on the IBM SAN Fibre Channel Switch. Figure 4-41 shows a switch connected to port 4 of the Gateway, and which has registered to the database. Since we cannot load any host registration software onto the switch, all other fields are left *unknown*.

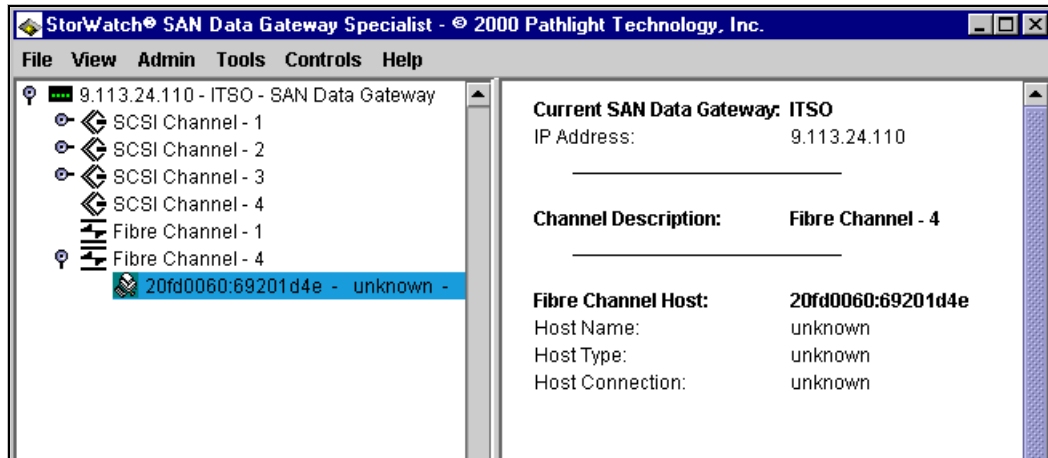


Figure 4-41 Switch registration

By selecting and highlighting the Fibre Channel port, as shown in Figure 4-42, you will see in the right-hand view pane that the port is now using a point to point, or N_port, connection that denotes a fabric login rather than a loop login.

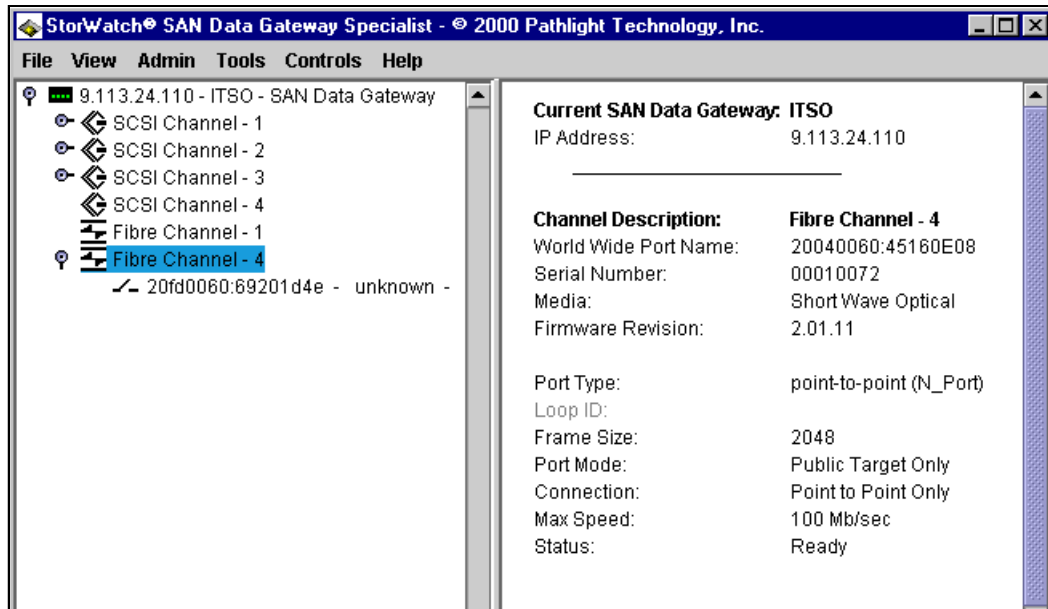


Figure 4-42 Switch port login

You can also check the port on the switch, by accessing the switch configuration from a Web browser, to ensure that the port has registered and is communicating properly. Figure 4-43 shows that, in this example, port 7 of the switch was used to connect to the Gateway. Note that the WWPN is the same as in the Specialist window and that the port type is F_port.




Figure 4-43 IBM SAN Fibre Channel Switch port settings

Changing the switch information

From the VPS Access Settings window, it is possible to change the *unknown* information of the switch. **Select Controls -> Virtual Private SAN**, and select the entry that has the WWN of the switch. You can now change the information to further describe the switch, or other pertinent information if desired. This is shown in Figure 4-44.

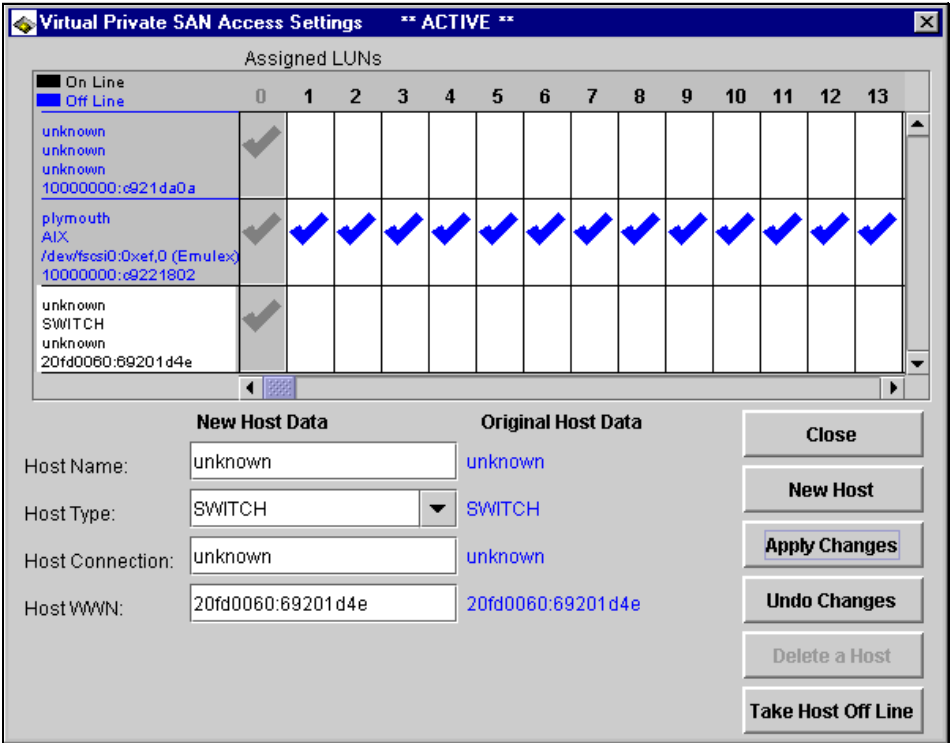


Figure 4-44 Changing switch information

The information on the host *plymouth*, that was attached using port 4, is still kept in the VPS database, but it is now shown in blue to indicate that it is offline.

Figure 4-45 is a diagram showing the configuration with the switch.




Figure 4-45 Adding an IBM switch to the Gateway

Once changed, as shown in Figure 4-46, the information shown on the main window will reflect the change, and the icon on the left-hand side of the WWN changes to depict a switch.

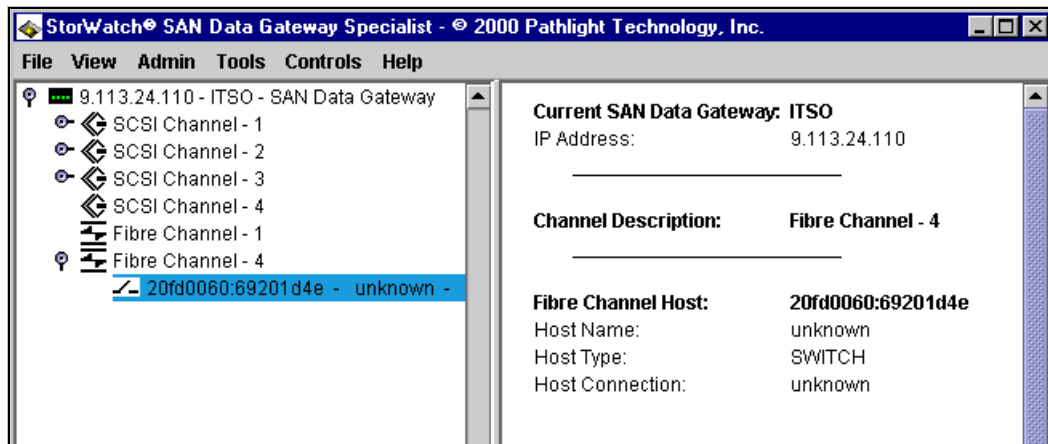


Figure 4-46 Switch port information

Attaching hosts to the switch

Any hosts that will attach to the switch should have the host registration software loaded and installed. Refer to , “Host registration” on page 298 for details. Plug in the Fibre Channel cable from their respective adapters and power on, or restart the host. The registration of the hosts will be completed through the switch. Figure 4-47 shows a switch installed on port 4 and two hosts connected on the switch. The host *plymouth* is now connected through the switch and an NT host was added.

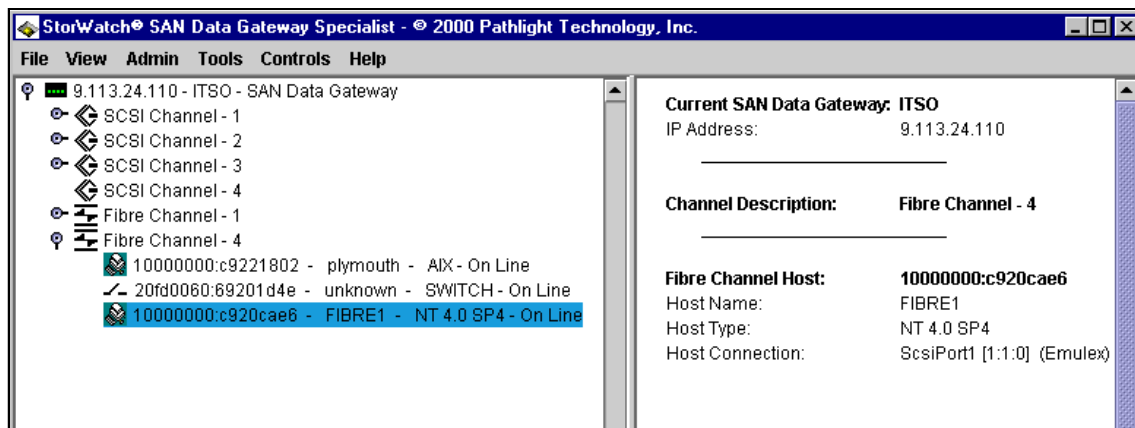


Figure 4-47 Hosts and a switch on a Gateway port

By selecting and highlighting the new NT host, we can see its information, which was automatically sent by the host registration process.

The configuration with two different hosts connected to the switch, as described previously, is shown in Figure 4-48.

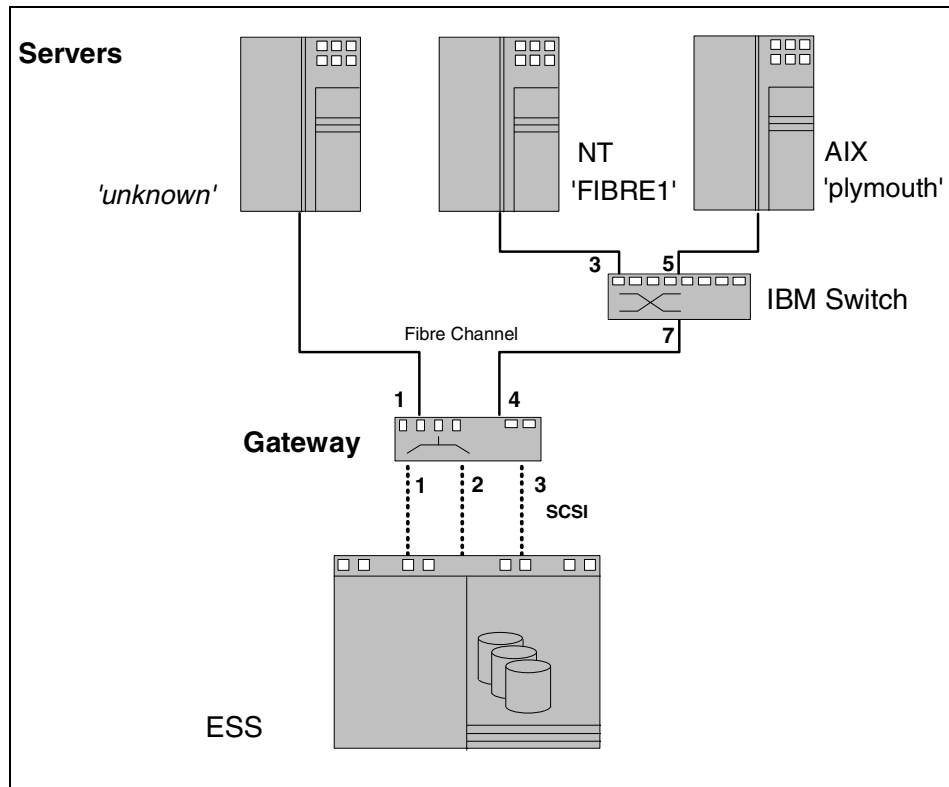


Figure 4-48 Adding two heterogeneous hosts to the switch

Device access

To have the new NT host access some of the devices, you need to set up the VPS Access parameters by selecting **Controls -> Access Options -> Virtual Private SAN**. Figure 4-49 shows the switch and the two hosts in the VPS database. As *plymouth* was previously connected direct to the Gateway, its settings have been maintained, but now that it is reconnected, it is back online. The NT host does not have any LUN access yet.

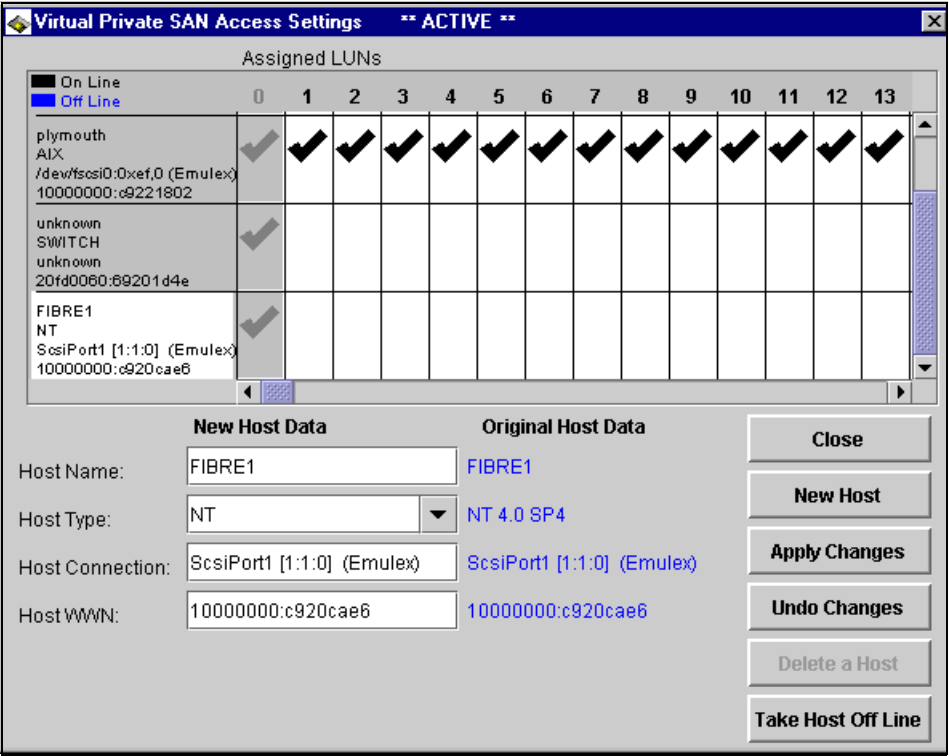


Figure 4-49 VPS Access window with switch and two hosts

A check mark in the box allows a host access to a LUN, or clear the box if you want to restrict access as described in , “Setting up a Virtual Private SAN” on page 303. Figure 4-50 shows the host *plymouth* with access to assigned LUNS 17 to 24, and the NT host *FIBRE1* is now set with access to LUNs 25 to 30.

The screenshot shows the 'Virtual Private SAN Access Settings' window with the title bar indicating it is '** ACTIVE **'. The window is divided into two main sections: 'Assigned LUNs' and 'New Host Data'.

Assigned LUNs: This section contains a table with columns for LUNs 17 through 30. The rows represent different hosts:

	17	18	19	20	21	22	23	24	25	26	27	28	29	30
plymouth AIX /dev/fscsi0:0x00:0 (Emulex) 10000000:c9221802	✓	✓	✓	✓	✓	✓	✓	✓						
unknown SWITCH unknown 20fd0060:69201d4e														
FIBRE1 NT ScsiPort1 [1:1:0] (Emulex) 10000000:c920cae6									✓	✓	✓	✓	✓	✓

New Host Data: This section contains input fields for 'Host Name:', 'Host Type:' (set to 'NT'), 'Host Connection:', and 'Host WWN:'.

Original Host Data: This section is currently empty.

Buttons: On the right side, there are buttons for 'Close', 'New Host', 'Apply Changes', 'Undo Changes', 'Delete a Host', and 'Take Host Off Line'.

Figure 4-50 Setting LUN access for the host FIBRE1

Once you select **Apply Changes** and then **Close**, the new settings will be in effect. For Windows NT to recognize the new devices, it will have to be restarted, so that it will do a rescan.

The *SWITCH* entry in the VPS database does not require any check marks to be set. Because the WWN of the host is known, and the switch WWPN is known, LUN access is accomplished through the switch as if it was transparent.

Note: LUN access is not only affected by the VPS and Zoning with the SAN Data Gateway, there is also Zoning within the IBM SAN Fibre Channel Switch. Be aware of the zone settings within the IBM SAN Fibre Channel Switch

ESS considerations

As you configure and allow access from host systems to volumes that are in the ESS through the Gateway, you must consider how the volumes were created within the ESS.

Volumes in the ESS are assigned to a particular SCSI port in the ESS. As they are assigned, you also specify the host type that will access these volumes.

Through the SAN Data Gateway it is possible to assign volumes to a specific operating system in the ESS, but then have a host with a different operating system access to these volumes.

As an example, look at Figure 4-48 on page 314. There we have a Windows NT and an AIX host. The ESS was configured and the Gateway attached so that SCSI 2 was assigned Windows NT volumes and SCSI 3 was assigned AIX volumes. It would be possible in the SAN Data Gateway to allow the NT host to access a LUN on SCSI 3 and the AIX host access to volumes on SCSI 2.

Here is where the device map would again be helpful in determining what Gateway assigned LUNs should be assigned to each host.

Volumes within the ESS assigned to a particular SCSI port should be of the same host type.

4.7 High availability considerations

Until now, our discussions of connections from the host to the SAN Data Gateway and from the Gateway to the ESS, or other SCSI devices, have been discussed with single host adapters or utilizing single SCSI connections. However, special considerations must be taken into account in deciding to add host adapters, or in providing redundant SCSI connections.

4.7.1 Single host with multiple Fibre Channel adapters

A host can have several adapters installed. If each adapter is assigned to different LUNs using the SAN Data Gateway, there are no contention issues. As each adapter is assigned specific LUNs, it cannot see or access any other LUNs that may be present.

However, there may be situations where the adapters are used to perform load balancing, and failover much like the SCSI connections from a host to the ESS that are in place today. As the host in the SCSI situation requires the Subsystem Device Driver (SDD) on each host, this is also true when using the Fibre Channel adapters.

Software, similar to SDD, is required in the host for it to recognize that it may have two or more paths to the same volume. If this was not loaded on the host, the host would recognize the volumes as being different, and there will be contention problems, and data corruption.

As SDD is currently only supported on SCSI adapters and not with Fibre Channel adapters, it is also not supported when using the SAN Data Gateway to connect to an ESS.

4.7.2 Multiple SCSI connections

The SAN Data Gateway can have all four SCSI channels connected to a single ESS. A volume in the ESS can be assigned to more than one SCSI port in the ESS. If these SCSI ports, that have the same volume assigned to it are connected to a SAN Data Gateway, the Gateway will assign multiple LUN ID numbers to the same volume. This is because the Gateway, upon discovery, or scanning of the SCSI buses, will view the volumes on each SCSI channel as separate volumes. For further explanation on this refer to 4.4.2, “Device discovery” on page 294.

If each LUN ID was then masked and zoned to different host systems, it is vital that the hosts have some access sharing software loaded to control access and avoid data corruption.

If the LUN IDs were assigned to the same host, then again software similar to SDD is required for the host to recognize that it has multiple paths to the same volume.

As stated earlier, SDD is not supported on the SAN Data Gateway connection to an ESS.

4.7.3 Adding Fibre Channel switches

As switches are added to allow for more host attachments, there are considerations about how many paths the host has to the SAN Data Gateway. If we refer back to Figure 4-48 on page 314, there is a single path from the switch to the Gateway.

If another Fibre Channel path from the switch to the Gateway was added, each host now has two paths to access the same LUNs. Each host will *see* the volumes twice. Once again, to prevent the host from recognizing the same volume twice, software similar to SDD is required.

Another option here is to utilize the zoning and LUN masking capabilities of the SAN Data Gateway. This would ensure that certain volumes can only be accessed on one Fibre Channel Gateway port and by a particular host.

Also available is to add zoning within the switch. The switch Fibre Channel ports can be zoned so that the host only has one path to the SAN Data Gateway. This would be used in combination with the zoning and LUN masking features of the SAN Data Gateway.



Implementing the Vicom Fibre Channel SLIC Router

The Vicom Fibre Channel SLIC Router, 7139-111, enables all IBM 7133, 7131, and 3527 SSA Serial Disk Systems to attach to host systems using Fibre Channel host adapters and drivers. This allows you to protect your investment in SSA disk, while being able to create and build a SAN infrastructure.

The Vicom Fibre Channel SLIC Router replicates data across or within serial disk systems — simultaneously mirroring two or three copies of data without host involvement. With global hot disk sparing, data is automatically rebuilt if a mirrored disk fails. In this way, the Vicom Fibre Channel SLIC Router improves performance and data availability while simplifying storage operations.

In the following sections, we cover these topics:

- ▶ “SLIC Router features” on page 322
- ▶ “Installing the SLIC Router” on page 323
- ▶ “SLIC Manager software” on page 326
- ▶ “Using SLIC Manager” on page 333
- ▶ “Composite drive” on page 337
- ▶ “Mirror drive” on page 342
- ▶ “Instant Copy drive” on page 347
- ▶ “Combining composite and mirroring” on page 352

5.1 SLIC Router features

The Instant Copy function can create a separately addressable copy of mirrored data that can be used for tape backup. After the backup has completed, data is resynchronized with the primary copy.

The Vicom Fibre Channel SLIC Router also can create composite drives by concatenating up to 16 physical disks.

Using these functions, physical drives become members of larger or more complex logical drives.

A diagram to depict a single host to Router configuration is shown in Figure 5-1.




Figure 5-1 SLIC Router with a single host

5.2 Installing the SLIC Router

To install and establish a SLIC storage system, the Router and all the disks to be used must be set up in a proper sequence. During this sequence, only *a single* SLIC Router must be used to do the configuration. Once configured, other SLIC Routers can be added to the loop. This sequence is described below.

Note: For information and a description to understand the LED codes that will be discussed, please refer to the *SLIC Router Installation and User's Guide*, 310-605759.

1. Power on SSA disks.

The SSA disks must be powered on to ensure that all disks spin up and pass the power sequence. Make sure all drive LEDs are on solid to indicate a functioning device. Any faulty or suspect disk drive should be replaced.

2. Clear Router Node Mapping.

To begin the Router installation, the first step is to clear the Node Mapping table. This is done by *shorting* the SSA ports on the rear of the Router. Plug in an SSA cable from port A1 to A2 on the back of the Router and power it on. Clearing the table will take only seconds, and when completed, the Status LED on the front of the Router will flash a code **060**. The Router is now powered off, the SSA shorting cable is removed and the SSA disks are attached.

3. Connect SSA disks to Router.

All the disks to be used are to be connected together to form a complete SSA loop with the Router included. All dip switches in SW1 should be in the down position. On SW2, dip switches 0 and 1 are set in the down position, all other switches in dip switch 2 should be turned up. This is considered *mode 3*. Power on the Router, the Status LED will begin to flash rapidly as it searches the SSA loop to recognize all the disk drives. This may take approximately 1 minute to complete. Once the Status LED has stopped flashing and is solid, the process is complete and the Router is powered down.

4. Perform Router Subsystem Diagnostic test.

A Subsystem Diagnostic test is now run on the disk drives called *mode 15*. This tests the disk drives for spindle spin up, read tests and nondestructive write tests. The Router is set for mode 15 by setting switches 0, 1, 2, and 3, on SW2, to the down position, and the rest turned up. The Router is now powered on, the Status LED will flash rapidly. The test will be done on each disk drive in the SSA loop separately and will begin with the drive closest to the A1 port on the back of the Router. As the test is completed on a drive, the

LED on the SSA drive will flash and then it will move to the next drive. This test should continue until all drives have been tested. The test runs in a continuous cycle, so once all drives have been tested at least once, the Router is powered off.

If a drive fails the test, the testing will stop, and the Router's Status LED will flash a diagnostic code. A code map with a description of the errors can be found in the *SLIC Router Installation and User's Guide*, 310-605759.

5. Assign Fibre Channel target.

With the Router powered off, you can now assign a Fibre Channel target ID number to the Router. Any number can be selected, however, this number must be a unique ID. No other device can have the same Fibre Channel target ID once it is set on the Router.

This is done by setting selected dip switches in SW1 to the down position. The switch is set up in binary notation: a switch that is down represents a 1 and a switch up represents a 0. Figure 5-2 shows the switch numbers and their corresponding value.

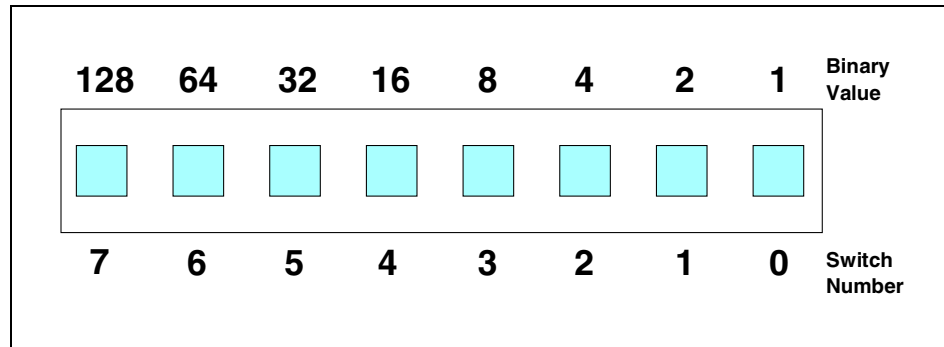


Figure 5-2 SW1 dip switches

For example, by setting switch 2 and 0 down, a value of 5 is represented. All other switches would be in the up position. By setting a switch down, its value is added. To change a number that was previously set on a Router, power must be cycled to the Router for the change to take effect.

6. Map the physical drives.

Before powering on again, SW2 must be set to mode 3 with switch 0 and 1 set down, and all other switches set up. The Router is powered on, the Status LED will flash rapidly to rediscover the disks and the SSA drive LEDs should be on solid. Once completed, the status LED will be solid, and the drives are now considered to be mapped physical drives. The Router is powered off.

7. Create general spares.

The mode on SW2 is changed to mode 12 to set all drives to be general spares. Mode 12 is represented by setting switch 2 and 3 down and the rest turned up. Power on the Router again, the Status LED will flash. After approximately one minute, the LED will flash code **100** to indicate it has completed. The Router is again powered off.

8. Format the drives.

The Router is set to mode 14, switch 1, 2, and 3 down on SW2, to format all disk drives. Power on the Router, depending on the number of drives, the format process will take anywhere from 30 to 60 minutes. During this time, the Status LED will flash rapidly and the SSA drive LEDs will flash. When it has completed, the Status LED will flash code **100**. Power off the Router.

9. Clear the node map.

Once completed, the node map *must* be cleared. This is done as described earlier by *shorting* ports A1 and A2 with an SSA cable. Power on, wait for code **060** to flash and then power off.

The drives are now ready to be assigned and used on a host.

You can also now setup mirror drives or composite drives within the Router. This is done by setting the switches in SW2 to other modes. For detailed information on setting the switches and selecting the other modes, please refer to the *SLIC Router Installation and Users Guide*, 310-605759.

10. Perform host attach and power up sequence.

For a host to now recognize and use the disks, set the dip switches in SW2 back to mode 3, this is normal host operation mode. The Fibre Channel cable from the host can be connected to the Router. If the SSA drives are not powered on, do this now, and this should be done before the Router. Next, the Router is powered on, wait for the Status LED to stop flashing and remain on solid. At this point the host can be powered on.

A check can be done to see that the SLIC Router is being recognized by the host adapter card. On a Windows NT with a QLogic Fibre Channel adapter, during boot up look for a prompt to enter the QLogic bios by entering in **ALT Q**. At the BIOS window, select **Scan Fibre Devices**. A list of the Fibre Channel target IDs are presented, scroll down to the ID that you set in SW1. You will see the WWN of the SLIC Router. Exit the bios and the system will reboot.

Once the system has started, you use a method to ensure that the host has access to all the drives. This is different depending on the operating system of the computer. For Windows NT, select **Start -> Programs -> Administrative Tools -> Disk Administrator**. This tool will report that new disks have been found and will be added to the system.

5.3 SLIC Manager software

Rather than using the dip switches to configure the features, another option is to use the SLIC Manager software. The SLIC Manager also provides configuration, monitoring and management capabilities of the SLIC Router and the SSA drive loop. The SLIC Manager can be setup to allow remote access if desired.

The Manager software consists of server and client portions. The server includes a daemon service and a user interface. The client has the user interface only.

The server portion must be loaded on the host that is directly attached to the Router, as the daemon service is started from here. The daemon must reside on the host that is directly connected to the Router. This host can also be used to run the Manager software for local access.

The client software can be loaded on to any computer, running a supported operating system, that can communicate to the host with the daemon service running. It must communicate to the server host using TCP/IP. This allows remote access to the Router and the storage loop. See Figure 5-3.

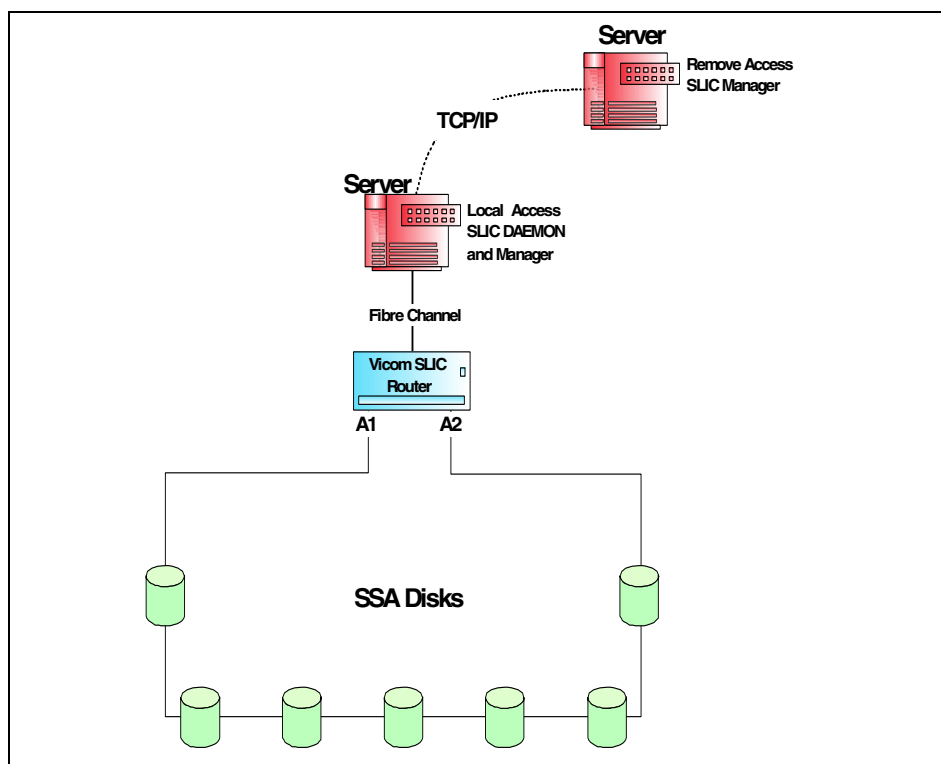


Figure 5-3 SLIC Manager access

The SLIC Manager has a graphical user interface (GUI) and a command line interface (CLI) available for Windows NT systems. UNIX systems will only have the command line interface available.

The following installation and configuration examples will detail using the GUI from a Windows NT platform. To review the commands that are available for UNIX platforms and installation instructions on other operating systems, please refer to the *SLIC Manager Installation and User Guide*, 310-605807.

5.3.1 Installing the SLIC Manager software

The SLIC Manager software can run on many operating systems. The following discussion will describe an installation on a Netfinity 5500 with Windows NT 4.0.

To install the SLIC Manager server software for local access, the Vicom Utilities CD-ROM is placed in the CD drive. Select **Start -> Run** and **Browse** the CD drive. Go to `slicmgr\i386\server\setup.exe` and click **OK**. Follow the prompts displayed on the window to install the Manager software. This will install the daemon service also.

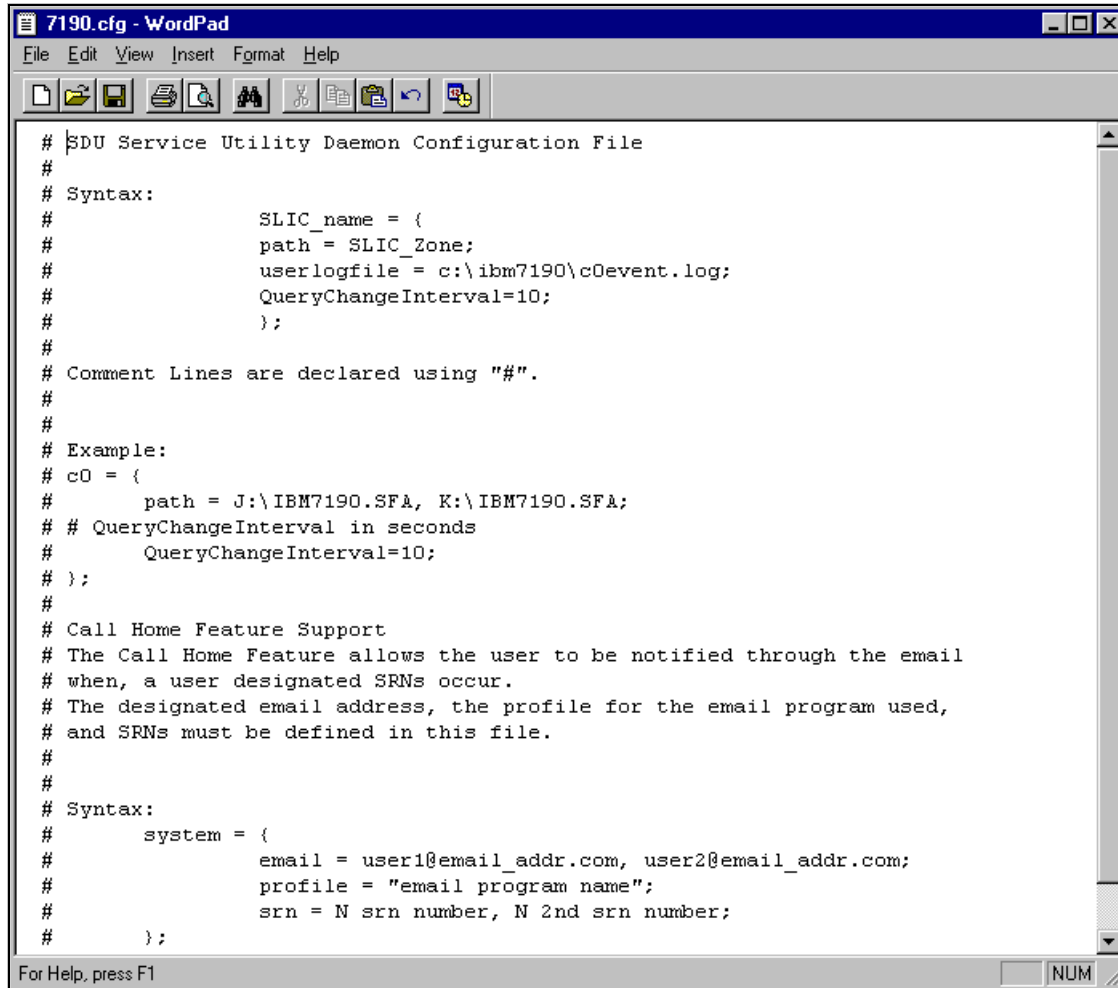
For remote or client access the daemon service is not required. To load the Manager software only, go to `slicmgr\i386\client\setup.exe` instead.

5.3.2 Communicating to the Router

For the SLIC Manager server software to communicate to the Router, it requires space on a disk or several disks that are within the SSA loop. This is referred to as a SLIC Zone. To create space on a disk, a file or partition — depending on the operating system used — is created for the Manager software to use. To create this SLIC Zone, a configuration file must be created or edited.

Editing the configuration file

When the software is loaded, a sample configuration file called *7190.cfg* is added in the *C:\ibm7190\sdus* directory. This is a text file that can be viewed and edited by simple text editors, such as Windows Wordpad. Open up the *7190.cfg* file and it will contain a sample of how the file should look. Also note that on the left hand side, the *#* sign is entered in every line to mark it out as a comment. This is shown in Figure 5-4.



```
# SDU Service Utility Daemon Configuration File
#
# Syntax:
#
#         SLIC_name = {
#             path = SLIC_Zone;
#             userlogfile = c:\ibm7190\c0event.log;
#             QueryChangeInterval=10;
#         };
#
# Comment Lines are declared using "#".
#
# Example:
# c0 = {
#     path = J:\IBM7190.SFA, K:\IBM7190.SFA;
#     # QueryChangeInterval in seconds
#     QueryChangeInterval=10;
# };
#
# Call Home Feature Support
# The Call Home Feature allows the user to be notified through the email
# when, a user designated SRNs occur.
# The designated email address, the profile for the email program used,
# and SRNs must be defined in this file.
#
# Syntax:
#
#     system = {
#         email = user1@email_addr.com, user2@email_addr.com;
#         profile = "email program name";
#         srn = N srn number, N 2nd srn number;
#     };
#
```

Figure 5-4 Sample configuration file

This file can now be edited to be used as the configuration file for your SLIC system. Begin by deleting the # sign on the lines that contain the sample configuration. The rest of the information can be entered in, as shown in Figure 5-5. A description of each entry field is also provided.

```
# SDU Service Utility Daemon Configuration File
#
# Syntax:
#
#         itso = {
#             path = F:\IBM7190.SFA;
#             userlogfile = c:\ibm7190\itsoevent.log;
#             QueryChangeInterval=10;
#         };
#
# Comment Lines are declared using "#".
#
#
# Example:
# c0 = {
#     path = J:\IBM7190.SFA, K:\IBM7190.SFA;
#     # QueryChangeInterval in seconds
#     QueryChangeInterval=10;
# };
#
# Call Home Feature Support
# The Call Home Feature allows the user to be notified through the email
# when, a user designated SRNs occur.
# The designated email address, the profile for the email program used,
# and SRNs must be defined in this file.
#
#
# Syntax:
#
#     system = {
#         email = user1@email_addr.com, user2@email_addr.com;
#         profile = "email program name";
#         srn = N srn number, N 2nd srn number;
#     };
#
```

For Help, press F1

Figure 5-5 Edited configuration file

Configuration file information

The *SLIC_name* can be any name that you would like to use to identify the Router.

Creating a SLIC Zone

The *path* refers to the SLIC Zone, file or partition, used for the Manager to communicate to the Router. To edit this option, it requires that a drive on the SSA loop has been recognized by the host and that the drive has been formatted. In the example above, a Windows NT host was used. The Windows NT Disk Administrator was used for the host to recognize the drives, the first drive assigned the next drive letter, F, and it was formatted.

The file naming for a SLIC Zone depends on the type of operating system running. For Windows NT, the naming is <drive letter>:\IBM7190.SFA. You can enter in many SLIC Zones, but only one is required to get access at the beginning. After the other drives have been configured as mirrors or composite drives, then SLIC Zones can be created for these drives if desired.

Including many SLIC zones in the path statement will allow the Manager to access a zone on another drive. This is helpful to protect against when a drive fails, and that drive has a SLIC zone defined to it. If the Manager cannot access the first SLIC zone, it would try the next zone in the order it was entered in the path statement.

For the naming conventions used on other operating systems to create a SLIC Zone, refer to the *SLIC Manager Installation and User Guide*, 310-605807.

The *userlogfile* will define a file with which you can view logged events.

The *QueryChangeInterval* sets the time in seconds that the daemon will poll the Router. The recommended time set here is 10.

Ensure that at the end of every line a semi-colon “;” is used, and that, if several SLIC Zones are created, a comma separates them. Save and exit the file.

Installing the SLIC Manager daemon

With the configuration file edited and a SLIC Zone created, the daemon service can be installed and run. To install the service in Windows NT, open a DOS prompt and go to C:\ibm7190\sdus. Type in **slicd -install**, and the daemon will be installed.

Starting the SLIC Manager daemon

To start the daemon service, select **Start -> Settings -> Control panel** from Windows NT. Double-click the **Services** icon. Scroll down until you see **Vicom SLIC Manager**; select and highlight it. You will see two columns to the right to indicate its status.

To start the service, click the **Start** button and it will take a few moments to complete. Once it is done, you will see the word **Started** in the Status column. If the Startup column contains the word **Automatic**, no further action is required. If not, click the **Startup** button, and change the Startup Type to **Automatic**. This will have the daemon service start automatically during a reboot. This is shown in Figure 5-6.

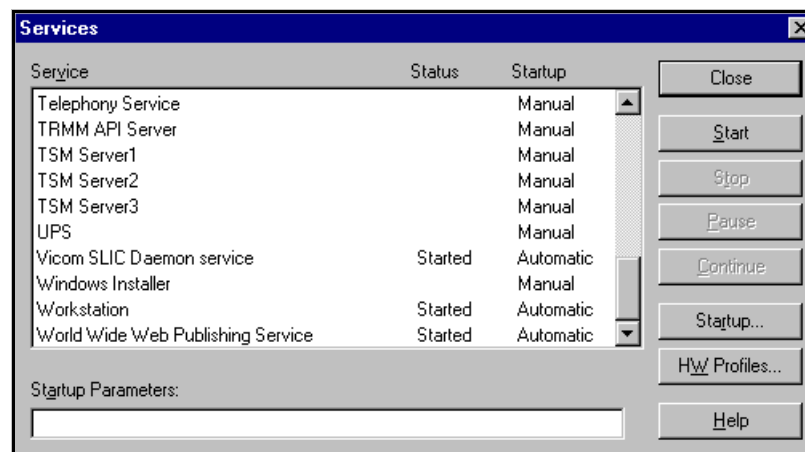


Figure 5-6 SLIC daemon start up in Windows NT

5.3.3 Starting the SLIC Manager

To start the Vicom SLIC Manager software, select **Start -> Programs -> Vicom -> Vicom SLIC Manager**. The software will load, and a dialog box will appear. In the box with the heading **Hostname**, enter in the name or IP address of the host the daemon service is running. Enter in the SLIC name you entered in when editing the *7190.cfg* file. An example is shown in Figure 5-7.




Figure 5-7 SLIC connection window

Click **OK** and the software will begin to communicate to the Router. You will notice that the top title bar of your window will now include the host name and SLIC name as in Figure 5-8.



Figure 5-8 SLIC Manager title bar

5.4 Using SLIC Manager

You can now look to see that all communications are working properly by going to the toolbar and selecting **Tools -> Control Center**. A dialog box will appear, as shown in Figure 5-9.

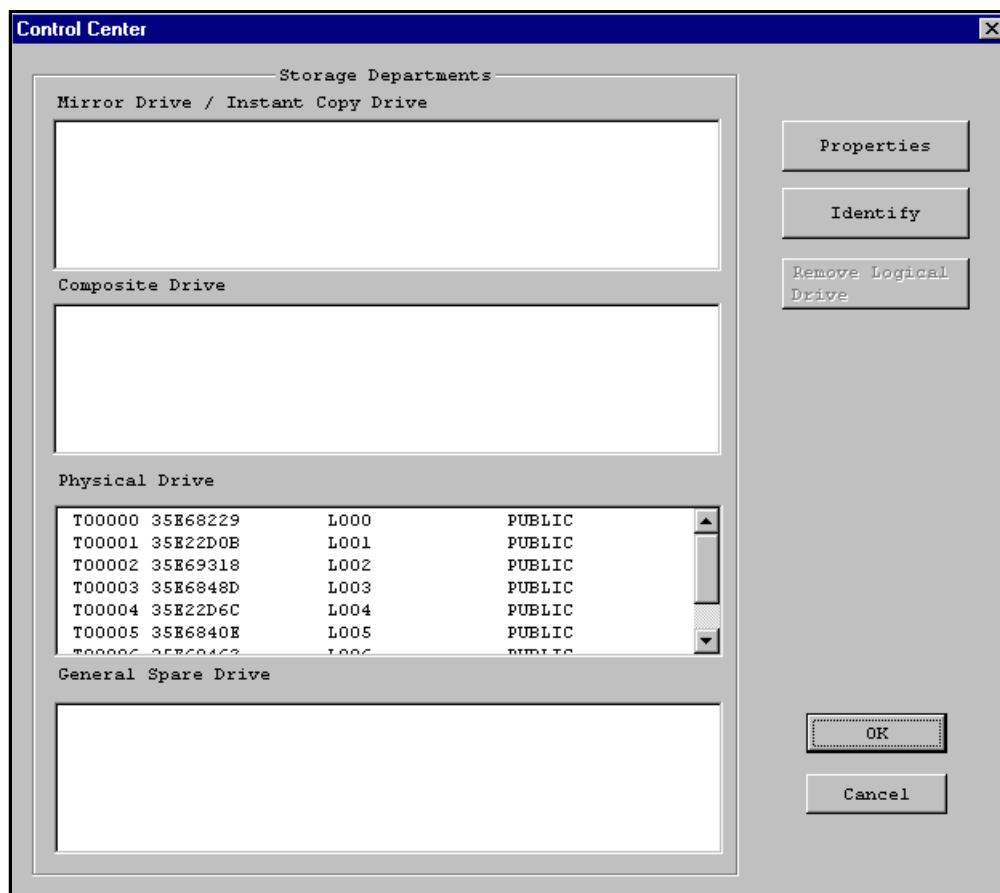


Figure 5-9 Control Center window

In the Physical Drive box, the drives that are on the SSA loop can be seen. This window will be useful as you start to create mirrors and composite drives, because it provides a summary of all drives.

5.4.1 Drive properties

You can get detailed information on each drive. Select the drive so that it is highlighted and then select **Properties**. A dialog box will appear with the drive's information, as shown in Figure 5-10.

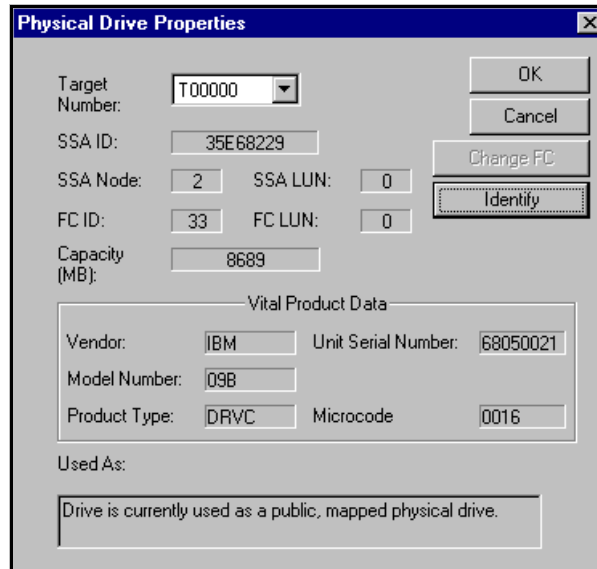


Figure 5-10 Disk drive properties

Here you can see its SSA attributes, its Fibre Channel attributes and its model type and serial number. By clicking the **Identify** button, the LED on the selected drive will begin to flash.

5.4.2 Router properties

To view the information on the Router, go to the toolbar and select **Properties -> SLIC Properties**. As shown in Figure 5-11, you will see the serial number of the Router, its ID that was set in SW1, and its supported features.

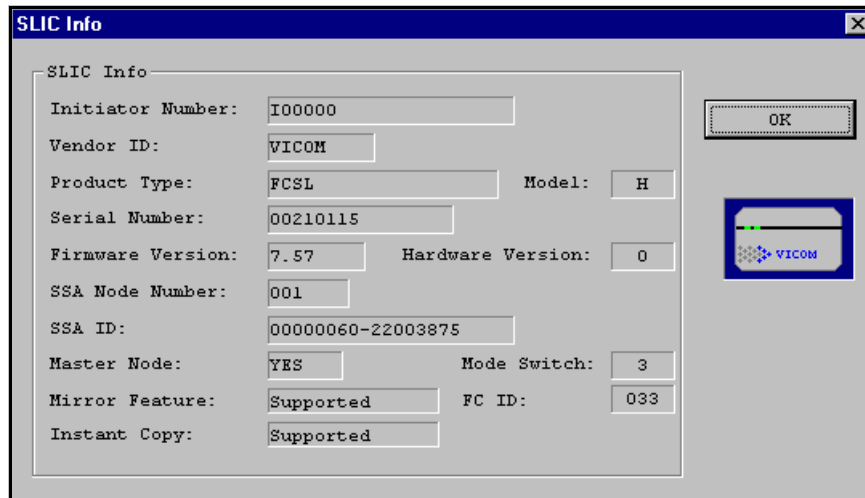


Figure 5-11 SLIC Router properties

5.4.3 Setting Router to master

As you move through the toolbar, you may notice that most selections have been grayed out. This is due to the fact that the Router is currently in a subordinate role and does not have access to create mirrors or composite drives. This function is done by a Master Router. There can be only one master in a SLIC loop. This is used as more Routers and more disks can be added to the loop. With several Routers in the same loop, there needs to be a requirement where one system acts as the control, and the others will follow and understand any configuration changes that may occur.

To set the Router into a master role, select **File -> Program Option** from the top toolbar. You will be presented a dialog box, as shown in Figure 5-12.




Figure 5-12 Setting the Router to master

Click in the **Always Master** box so that a check mark appears. Once you click **OK**, the Router will then be set as the master. You will notice that now all options in the toolbar are available and can start to use the features of the SLIC Router.

By placing a check mark in the Auto Start box, the SLIC Manager will automatically connect to the Router defined in the Connection window, as seen in Figure 5-7 on page 332.

5.4.4 The SignOn drive

When the SLIC zone was created to be used as the communication path, a disk file or partition was created on a specific disk within the SSA loop. As you begin to access the features of the SLIC Router, it should be known which disk was used to create the SLIC zone. This disk is considered to be the SignOn drive.

In the topics 5.5, “Composite drive” on page 337 and 5.6, “Mirror drive” on page 342 we describe creating composite and mirror drives, and you will see that the properties of the individual physical drives may change. As they become part of a logical drive, they take on the properties of this logical drive.

If the SignOn drive is used to create a logical drive, its attributes may change and you may lose the communication path that was created in the SLIC zone. When you select the **SignOn** drive as a member of a logical drive, a dialog box will be displayed as in Figure 5-13 to remind you that the attributes of this drive may be affected.




Figure 5-13 Selecting SignOn drive dialog box

As long as the LUN number of the SignOn drive becomes the LUN of the new logical drive, the communications from the Manager to the Router will not be affected.

Another way to be certain that you do not lose your SignOn drive is not to use the SignOn drive to create logical drives. However, once some logical drives have been created, a SLIC zone can be created to one or more of the newly created logical drives. This logical drive can now be used as the SignOn drive and the previous drive is now available to be configured without any problems.

5.5 Composite drive

A composite drive is a large drive that consists of two or more smaller drives. The capacity of the composite drive is an aggregate of the capacities of all the smaller drives that are used to comprise this one large drive.

5.5.1 Creating a composite drive

To create a composite drive from the SLIC Manager, select **Tools -> Composite Drive Setup Wizard**. A dialog box, Composite Drive List, will appear. Currently, the list will be blank, because there are no composite drives created. Once there are composite drives created, you will see a list of the drives. Click the **Next** button and you will see the Members Selection window, as shown in Figure 5-14.

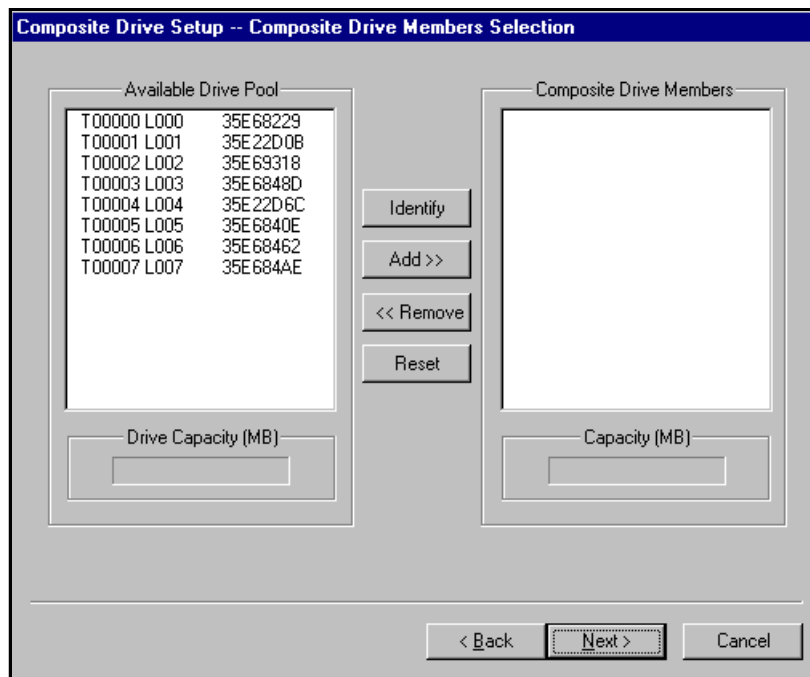


Figure 5-14 Composite Drive Member Selection window

From the Available Drive Pool list, click a desired drive and then click the **Add>>** button. The drive name will be added to the Member window. An asterisk will appear on the left hand side of the drive that was selected in the Available Drive window, to denote that the drive has been selected. Each drive is added one at a time. To remove a drive from the Member window, select the desired drive and click the **Remove<<** button.

Below each window there is a *Drive Capacity* box. As a drive is selected, its capacity in megabytes is shown. As you add more member drives to the Member window, the *Drive Capacity* box will add all drive sizes together to provide a total capacity in megabytes. This is shown in Figure 5-15.

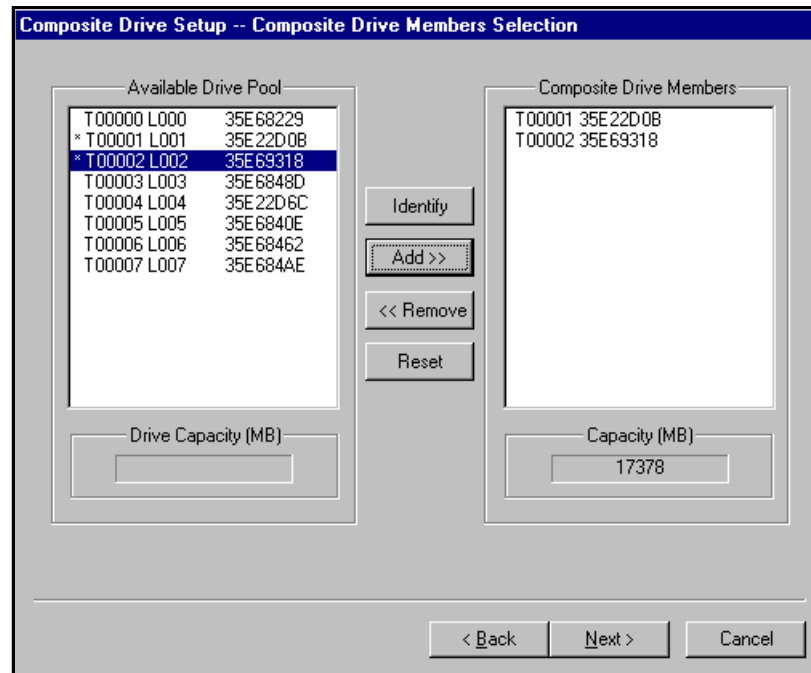
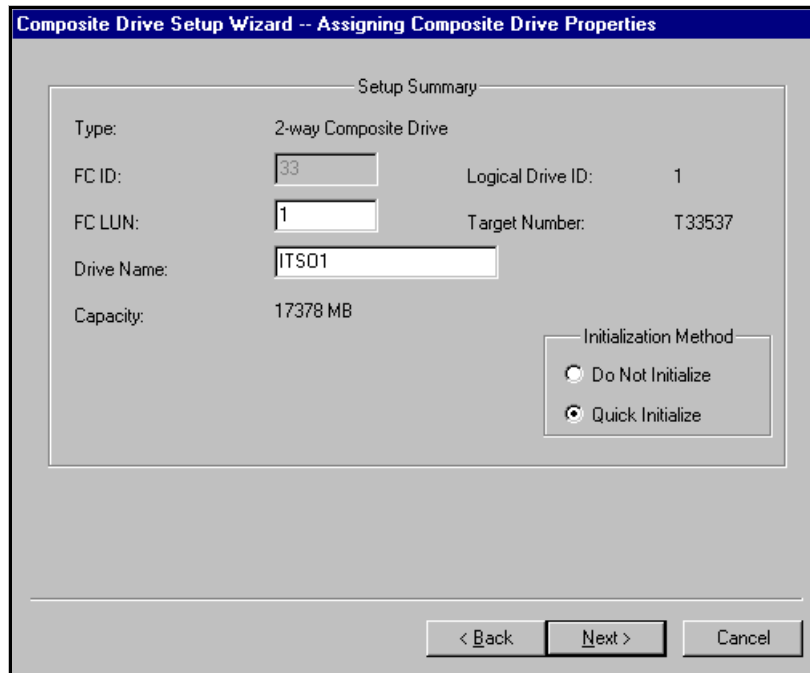


Figure 5-15 Creating composite drive from available drives

When all the desired drives are added, click **Next>**. The Assigning Properties window opens, as shown in Figure 5-16.



The image shows a Windows-style dialog box titled "Composite Drive Setup Wizard -- Assigning Composite Drive Properties". The dialog has a "Setup Summary" section with the following fields and values:

Property	Value
Type:	2-way Composite Drive
FC ID:	33
Logical Drive ID:	1
FC LUN:	1
Target Number:	T33537
Drive Name:	ITS01
Capacity:	17378 MB

Below the "Setup Summary" section is an "Initialization Method" group box containing two radio buttons:

- ☐ Do Not Initialize
- ☒ Quick Initialize

At the bottom of the dialog are three buttons: "< Back", "Next >", and "Cancel".

Figure 5-16 Assigning Composite Drive Properties window

The properties that can be changed are the FC LUN and the Drive Name. There will be a suggested LUN number in this field that can be accepted. If not, simply type in the desired LUN number. The name can also be defined to the composite drive for easier identification, with a limit of up to eight characters.

The Initialization Method box refers to whether or not to allow the operating system to write its signature on the composite drive.

Click the **Next>** button and a dialog box will appear, as shown in Figure 5-17, to allow you to create another composite drive. Click **Yes** if you would like to create another composite drive, and the Composite Drive List window opens and the steps described above can be repeated.

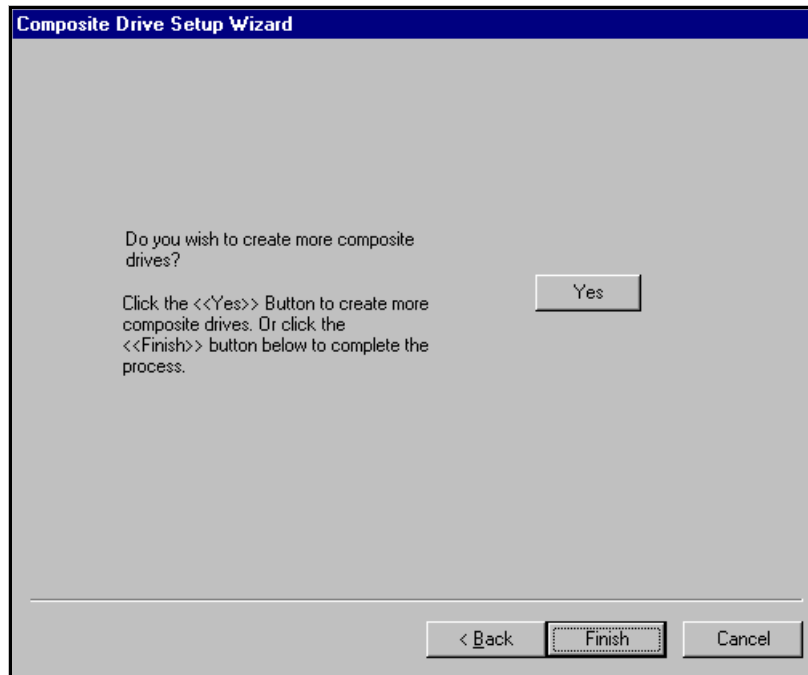


Figure 5-17 Completing the Composite Drive setup

Click **Finish** when you have created all the desired composite drives. Up to this point, the configuration has been kept within the SLIC Manager software. When the **Finish** button is clicked, the SLIC Manager will now communicate to the Router to complete the process and update the Router to control the drives.

The Host system must re-scan for devices, or restart, to be able to see the composite drive.

5.5.2 Composite drive properties

If you view the Control Center again, by selecting **Tools -> Control Center**, as shown in Figure 5-18, the newly created composite drive is listed in the Composite Drive box.

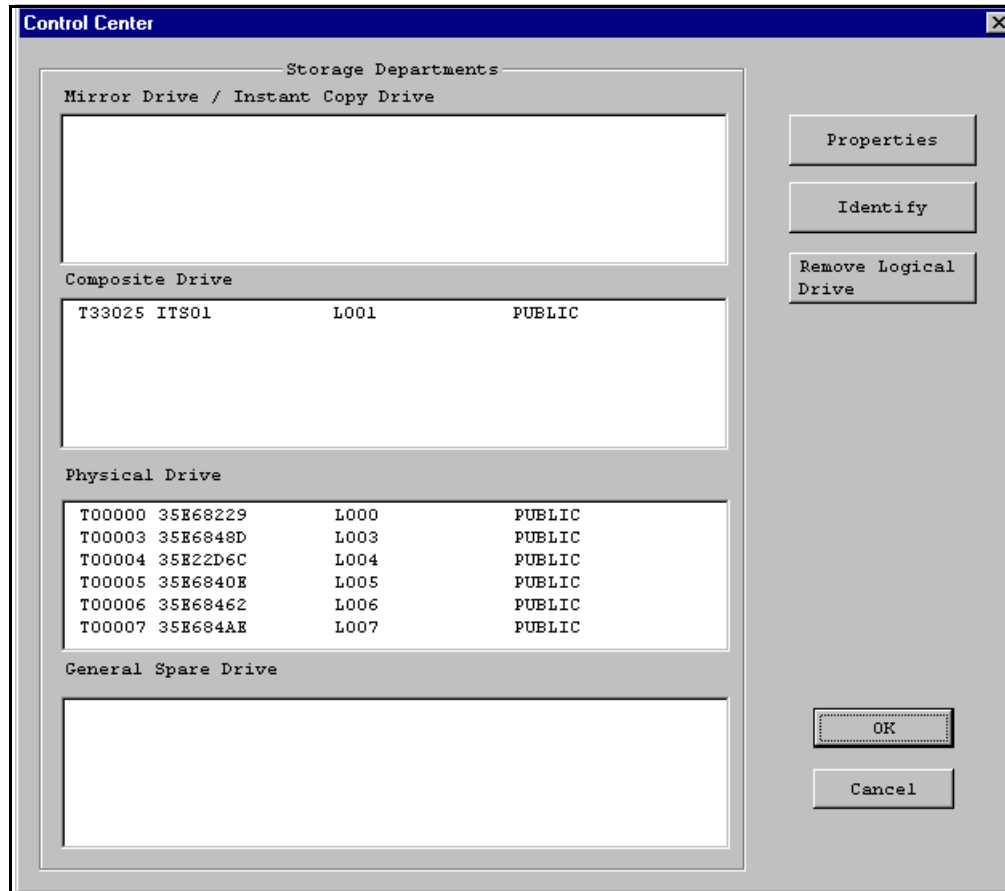


Figure 5-18 Control Center with composite drive

Select the composite drive and then click the **Properties** button. The Composite Drive Properties dialog box opens, as shown in Figure 5-19.

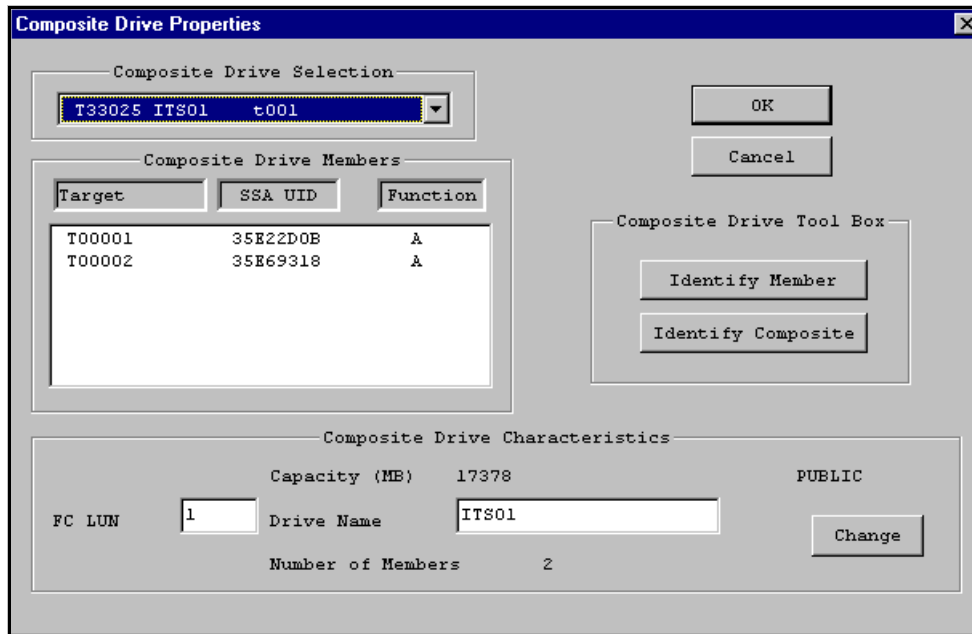


Figure 5-19 Composite Drive Properties

Here you can find information on the Composite Drive about the member drives that make up the composite, as well as the Composite Drive characteristics. Clicking on the **Identify Composite** button will cause the LED on the actual SSA drives, that belong to the Composite, to flash. If a member drive is selected and the **Identify Member** button is chosen, then the LED only on that drive will flash. In both cases a dialog box will appear to allow you to stop the flashing.

5.6 Mirror drive

A 2-way mirror drive has two drives that contain exactly the same information. The SLIC Router can also support a 3-way mirror or 1-way mirror. A 3-way mirror consists of three drives with the same information. A 1-way mirror is a single drive, or single composite drive, that is used with an Instant Copy Drive that can attach to the single drive mirror to synchronize the data. The Instant Copy Drive can then be split off from the mirror to perform a backup or other action.

The Instant Copy Drive feature can be used with 2-way and 3-way mirrors as well.

5.6.1 Creating a mirror drive

To create a mirror using physical drives, from the toolbar, select **Tools -> Mirror Drive Setup Wizard**. You will see a dialog box, Mirror Drive List, that will be blank. If there were mirror drives created, then it would display the names of the drives. Click the **Next>>** button and the **Mirror Drive Members Selection** window opens. The window on the left named *Available Drive Pool* contains a list off all drives that are candidates to participate in a mirror drive.

Select a drive by highlighting it and click the **Add>>** button. The drive name will be added to the Member window. An asterisk will appear on the left hand side of the drive just selected in the Available Drive window, to denote that the drive has been selected. A second or third drive can be added to create a 2-way, or 3-way mirror, respectively. Each drive is added one at a time. To remove a drive from the Member window, select the desired drive and click the **Remove<<** button.

An example of adding two drives to create a 2-way mirror is shown in Figure 5-20.

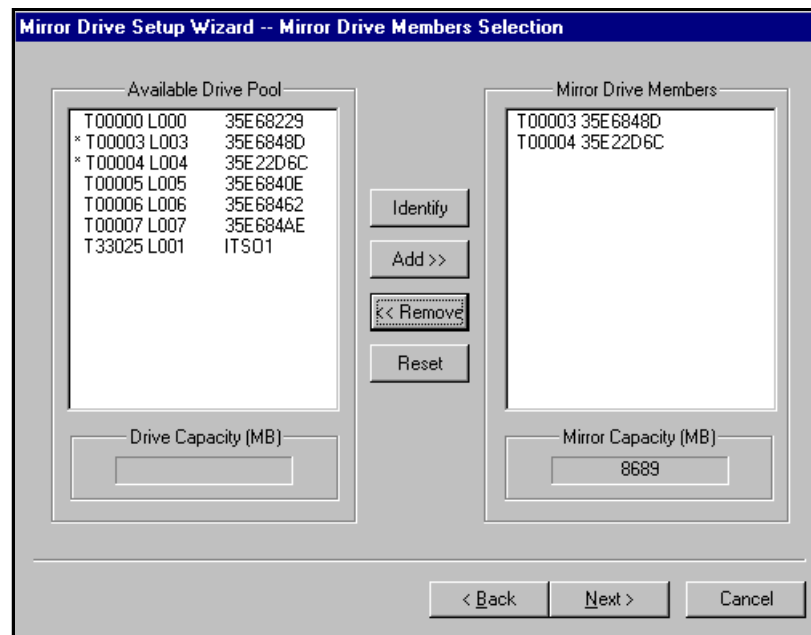


Figure 5-20 Mirror drive member selection

Below each window there is a Capacity window that will display the size of the available drive, or of the mirror drive. Each drive that participates in a mirror should be of equal capacity. If you select a drive that has a greater capacity and add it to the mirror, the mirror capacity will still be the smaller of the two, and the rest of the capacity of the larger drive will be unused. For example, if you added a 18 GB drive to the mirror in Figure 5-20, the Mirror Capacity window would still show the capacity of 8,696 MB. Approximately half of the 18 GB drive will be unused.

After all drives have been added, click **Next>** and you will be able to add a dedicated spare drive to the mirror if desired. Highlight one of the remaining available drives, click **Add>>** and its name will appear in the Mirror Drive Dedicated Spare window, as shown in Figure 5-21.

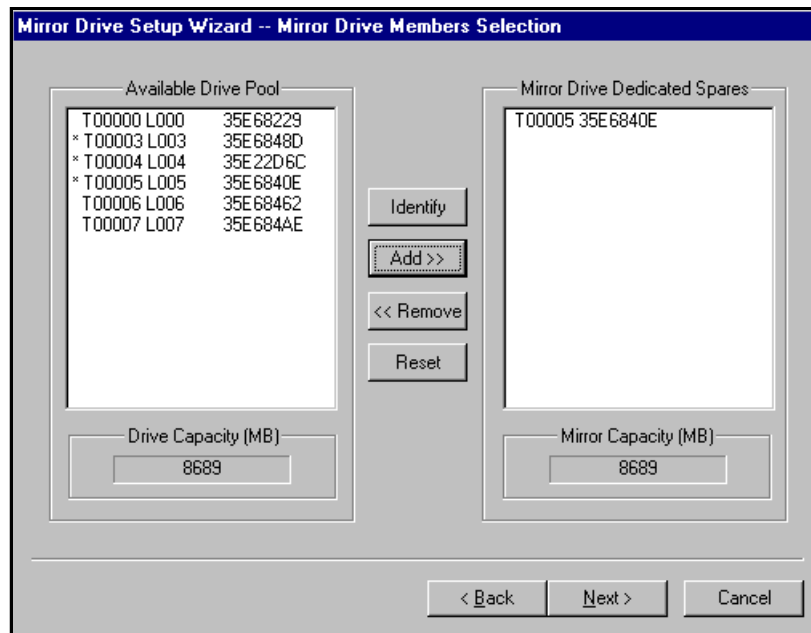
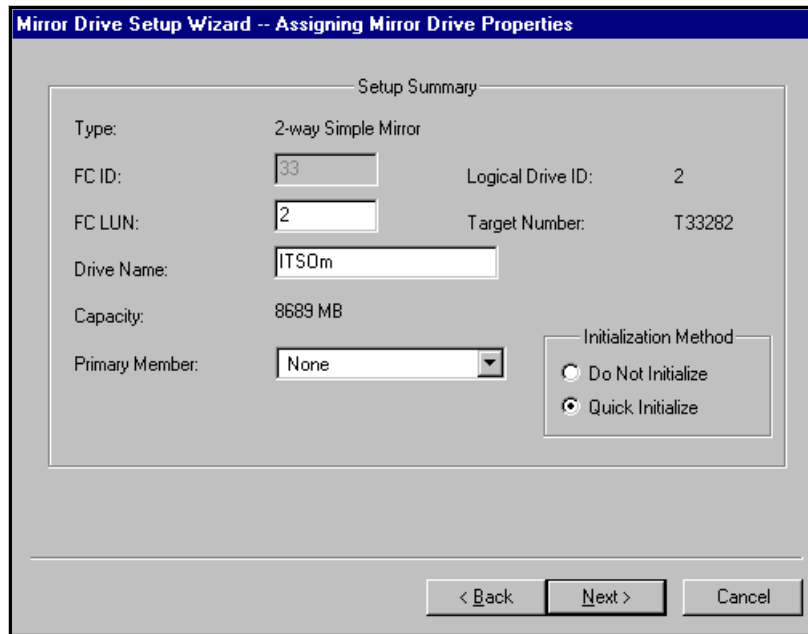


Figure 5-21 Adding a dedicated spare

Click the **Next>** button and the properties of the mirror drive can be changed. The properties that can be changed are the FC LUN and the Drive Name. There will be a suggested LUN number in this field that can be accepted. If not, simply type in the desired LUN number. A name can also be defined to the mirror drive for easier identification, with a limit of up to eight characters.

The Initialization Method box refers to whether or not to allow the operating system to write its signature on the Mirror drive.

The Assigning Mirror Drive Properties window is shown in Figure 5-22.



The image shows a Windows-style dialog box titled "Mirror Drive Setup Wizard -- Assigning Mirror Drive Properties". The dialog has a "Setup Summary" section containing the following fields and values:

Field	Value
Type:	2-way Simple Mirror
FC ID:	33
Logical Drive ID:	2
FC LUN:	2
Target Number:	T33282
Drive Name:	ITS0m
Capacity:	8689 MB
Primary Member:	None

Below the "Primary Member" field is an "Initialization Method" section with two radio buttons:

- ☐ Do Not Initialize
- ☒ Quick Initialize

At the bottom of the dialog are three buttons: "< Back", "Next >", and "Cancel".

Figure 5-22 Mirror drive properties

Click the **Next>** button and a dialog box appears to allow you to create another mirror drive. Click **Yes**, if you would like to create another mirror drive, and the Mirror Drive List window opens, and the steps described above can be repeated.

Click **Finish** when you have created all the desired mirror drives. Up to this point, the configuration has been kept within the SLIC Manager software. When the **Finish** button is clicked, the SLIC Manager will now communicate to the Router to complete the process and update the Router to control the drives.

If **Quick Initialize** in the Initialization Method box was selected, the Router will take a short period of time to write the host signature and build the mirror. During this time if you try to communicate to the Router, you may experience a slower than normal response.

5.6.2 Mirror drive properties

If you go to the Control Center window by selecting **Tools -> Control Center**, you will see that the mirror drive is now displayed in the Mirror Drive window. This is shown in Figure 5-23.

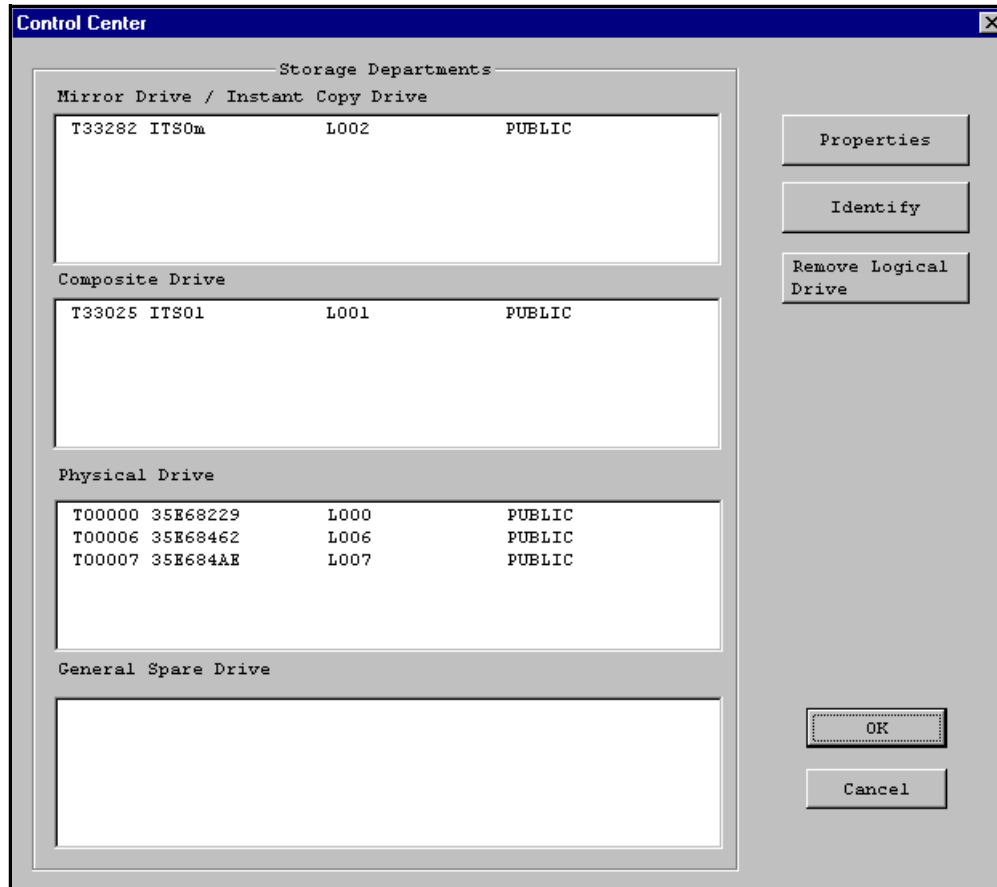


Figure 5-23 Control Center with Mirror Drive

If you select and highlight the mirror drive and then click the **Properties** button, the Mirror Properties window opens, and you can see the information in the mirror drive. Figure 5-24 shows an example of the properties of the mirror drive.

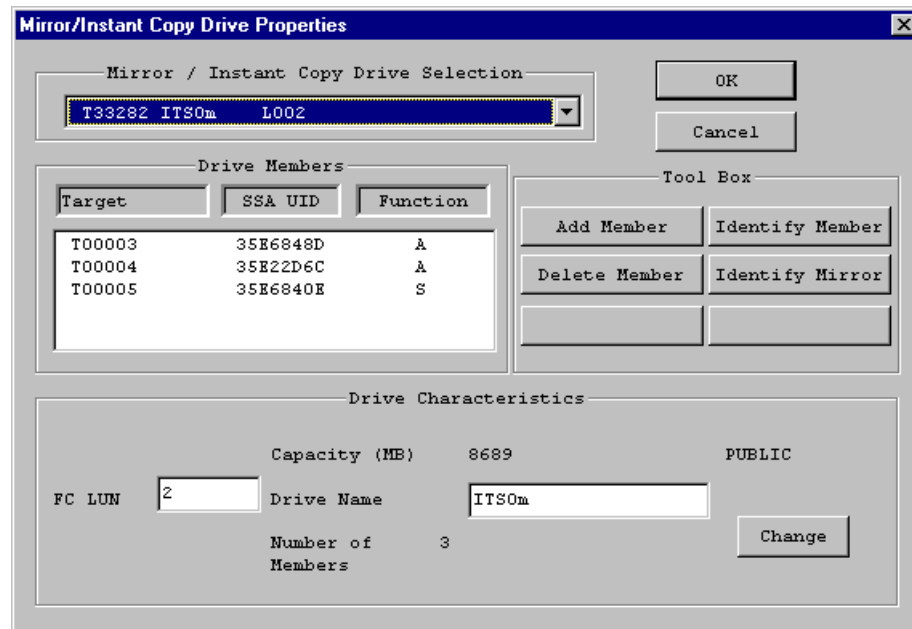


Figure 5-24 Mirror Drive Properties

Clicking on the **Identify Mirror** button will cause the LED on the actual SSA drives that belong to the mirror, to flash. If a member drive is selected and the **Identify Member** button is chosen, then the LED only on that drive will flash. In both cases, a dialog box appears to allow you to stop the flashing.

5.7 Instant Copy drive

Instant Copy is a feature that allows a drive to become part of a mirror, synchronize to the latest data, and then detach from the mirror. The drive can then be used to back up the data or used elsewhere if desired.

5.7.1 Creating an Instant Copy drive

To create an Instant Copy drive, select **Tools -> Instant Copy Drive Setup Wizard**. You will see a dialog box, Instant Copy Drive List, that will be blank. If there were copy drives created, it would display the names of the drives. Click the **Next>>** button and the **Instant Copy Drive Members Selection** window is displayed. The window on the left named *Available Drive Pool* contains a list off all drives that are candidates to become a copy drive.

Select a drive by highlighting it and click the **Add>>** button. The drive name will be added to the Member window. An asterisk will appear on the left hand side of the drive just selected in the Available Drive window to denote that the drive has been selected. An example is shown in Figure 5-25.

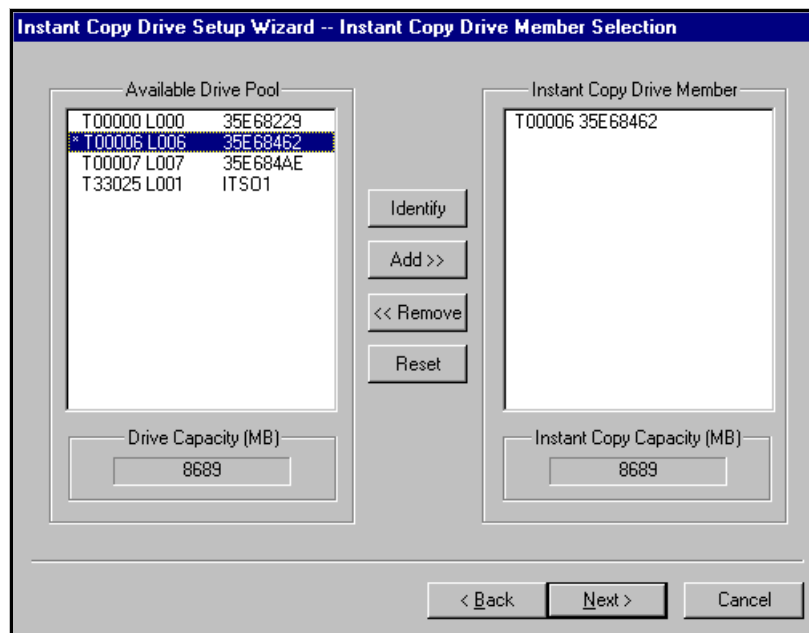


Figure 5-25 Instant Copy Drive Member Selection

To remove a drive from the Member window, select the desired drive and click the **Remove<<** button. Below each window there is a Drive Capacity box. As a drive is selected, its capacity in megabytes is shown. Click the **Next>** button to continue to the Assigning Instant Copy Drive Properties window, as shown in Figure 5-26.

Instant Copy Drive Setup Wizard -- Assigning Instant Copy Drive Properties

Setup Summary

Type:	Instant Copy Drive with Simple Member		
FC ID:	<input type="text" value="33"/>	Logical Drive ID:	3
FC LUN:	<input type="text" value="3"/>	Target Number:	T33283
Drive Name:	<input type="text" value="copy1"/>		
Capacity:	8689 MB		
Primary Member:	<input type="text"/>	Initialization Method <input type="radio"/> Do Not Initialize <input type="radio"/> Quick Initialize	

< Back Next > Cancel

Figure 5-26 Instant Copy Drive Properties

The properties that can be changed are the FC LUN and the Drive Name. There will be a suggested LUN number in this field that can be accepted. If not, simply type in the desired LUN number. A name can also be defined to the copy drive for easier identification, with a limit of up to eight characters.

Click the **Next>** button and a dialog box appears to allow you to create another copy drive. Click **Yes** if you would like to create another copy drive and the Instant Copy Drive List window opens, and the steps described above can be repeated.

Click **Finish** when you have created all the desired copy drives. Up to this point, the configuration has been kept within the SLIC Manager software. When the **Finish** button is clicked, the SLIC Manager will now communicate to the Router to complete the process and update the Router to control the drives.

5.7.2 Instant copy drive properties

You can go to the Control Center window by selecting **Tools -> Control Center**. The copy drive that was created above can now be seen in the Mirror Drive/Instant Copy window, as shown in Figure 5-27.

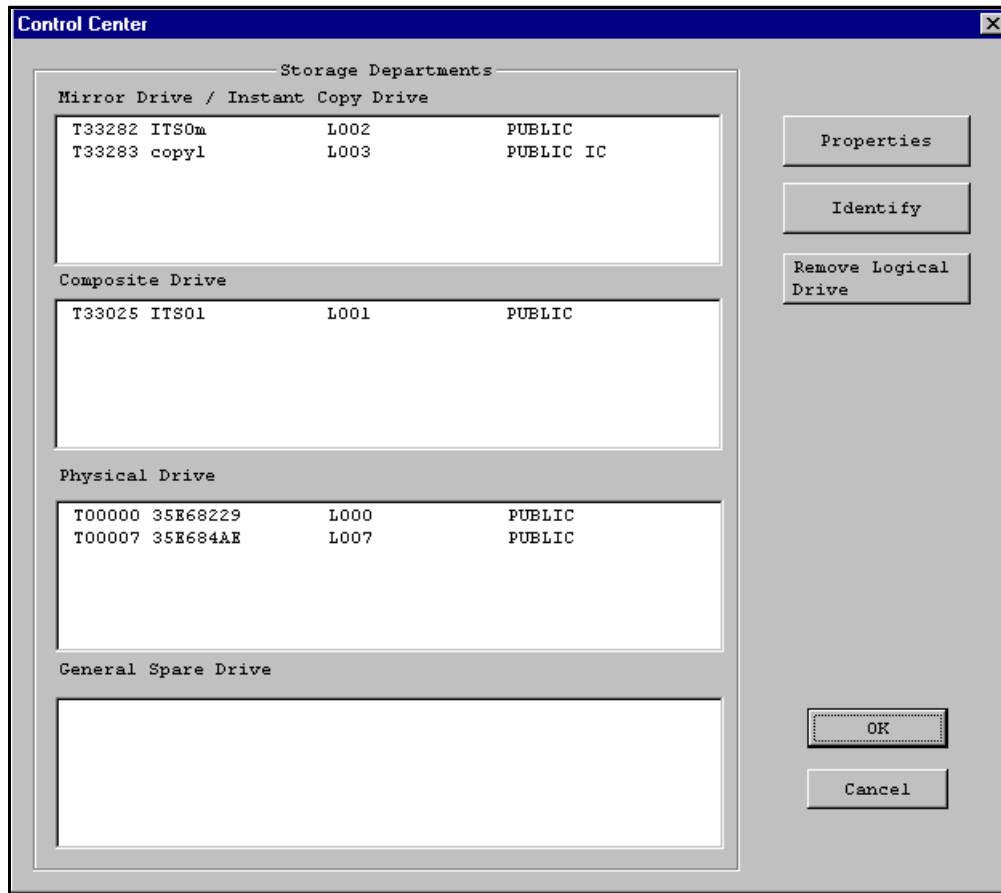


Figure 5-27 Control Center with Instant Copy Drive

Notice that in the information provided for the copy drive, there is an **IC** included to distinguish between mirror drives and copy drives within this window.

5.7.3 Adding an Instant Copy Drive to a mirror

To add or detach the copy drive from a mirror, you select and highlight the mirror drive, and then click the **Properties** button. The Mirror Drive Properties window opens, as shown in Figure 5-24 on page 347. Click the **Add Member** button and the Add Mirror Member window opens, as shown in Figure 5-28.

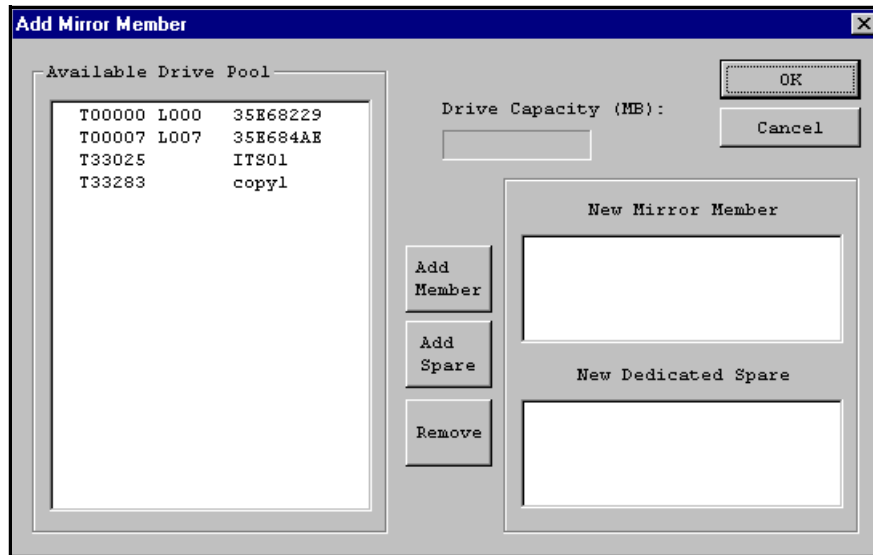


Figure 5-28 Add Mirror Member display

Select and highlight the copy drive from the Available Drive Pool window, click the **Add Member** button, and the name of the copy drive will appear in the New Mirror Member window. This is shown in Figure 5-29.

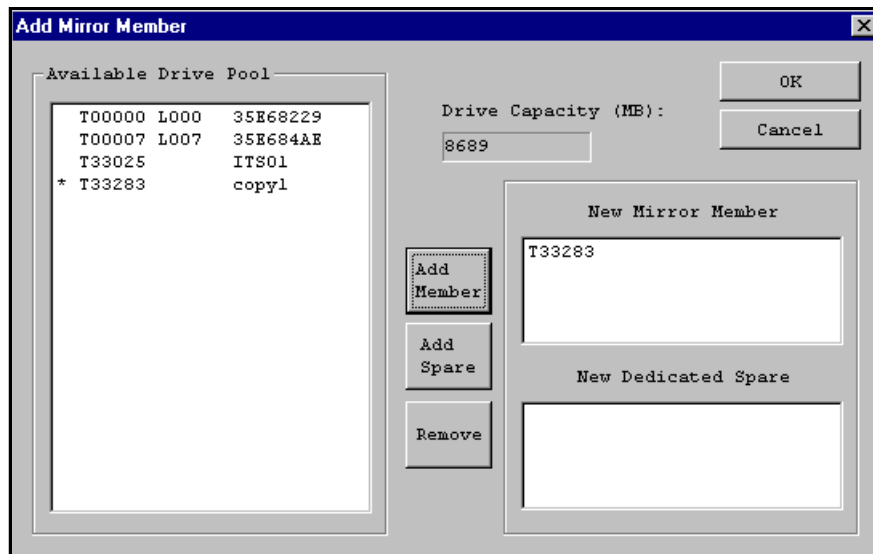


Figure 5-29 Adding drive members to a mirror

Click the **OK** button, and the Mirror Drive Properties will now reflect the change, as shown in Figure 5-30.

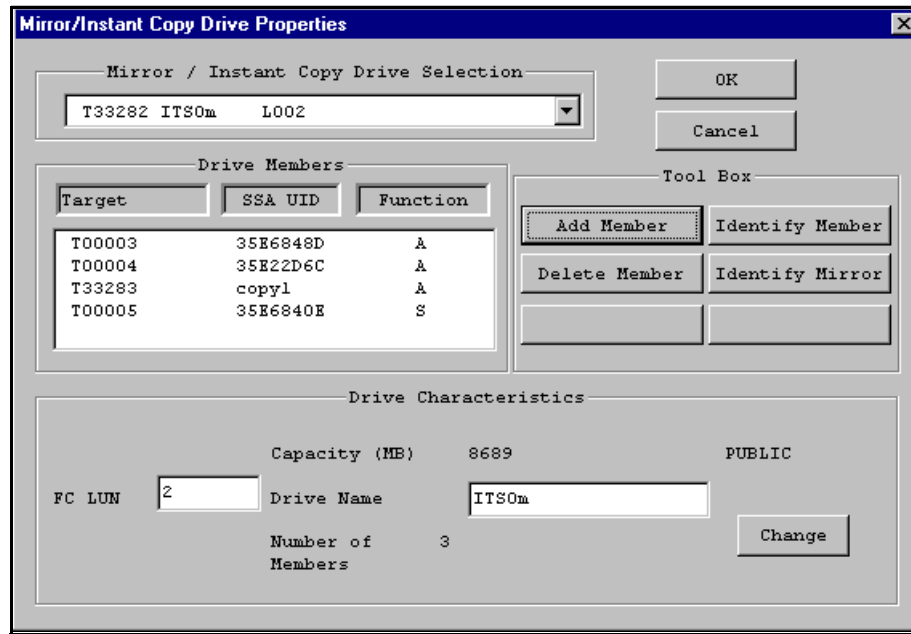


Figure 5-30 Mirror drive properties with copy drive attached

Click **OK** to complete the process.

5.7.4 Detach Instant Copy Drive from a mirror

To detach, or split off the copy drive from the mirror, the procedure is similar except at the Mirror Drive Properties window, select **Delete Member**. A window will appear that displays all current members of the Mirror. Select the Copy drive, and then delete it from the Mirror. The Copy drive can now be accessed by another host.

5.8 Combining composite and mirroring

The SLIC Manager can also be used to combine the two features of the Router. You can create a mirror drive using composite drives. A mirror can have drive members of different sizes, but the actual mirror capacity will be the smaller of the drive sizes.

5.8.1 Creating a second composite drive

To provide an example of a mirror using only composite drives, another composite drive is required. The example shown in Figure 5-31 shows that drive 6 and 7 were used to create another composite drive.

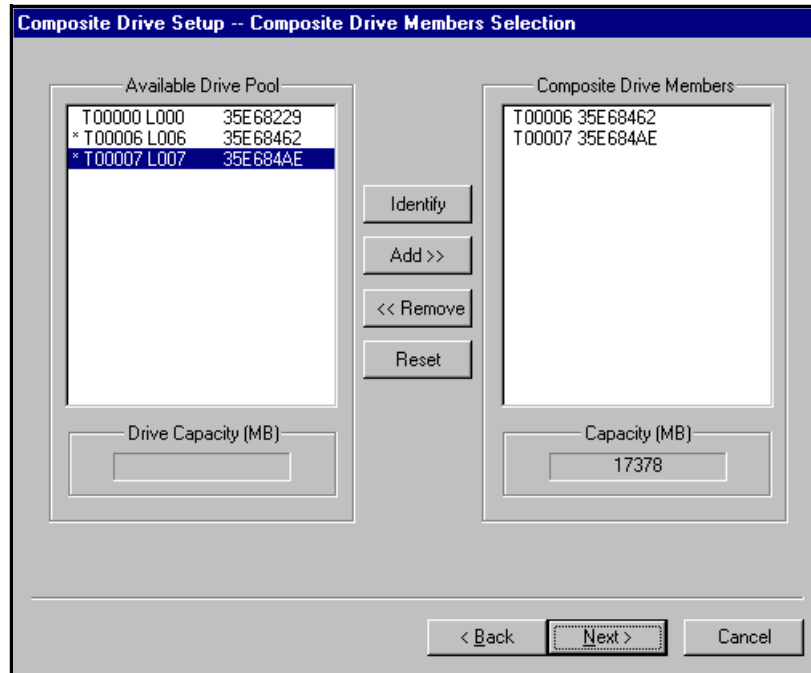


Figure 5-31 Creating composite drive to be used in a mirror

Follow the steps to create a composite drive as described in 5.5, “Composite drive” on page 337. Once created, you can view the Control Center window by selecting **Tools -> Control Center** from the toolbar. Figure 5-32 shows that there are now two composite drives.

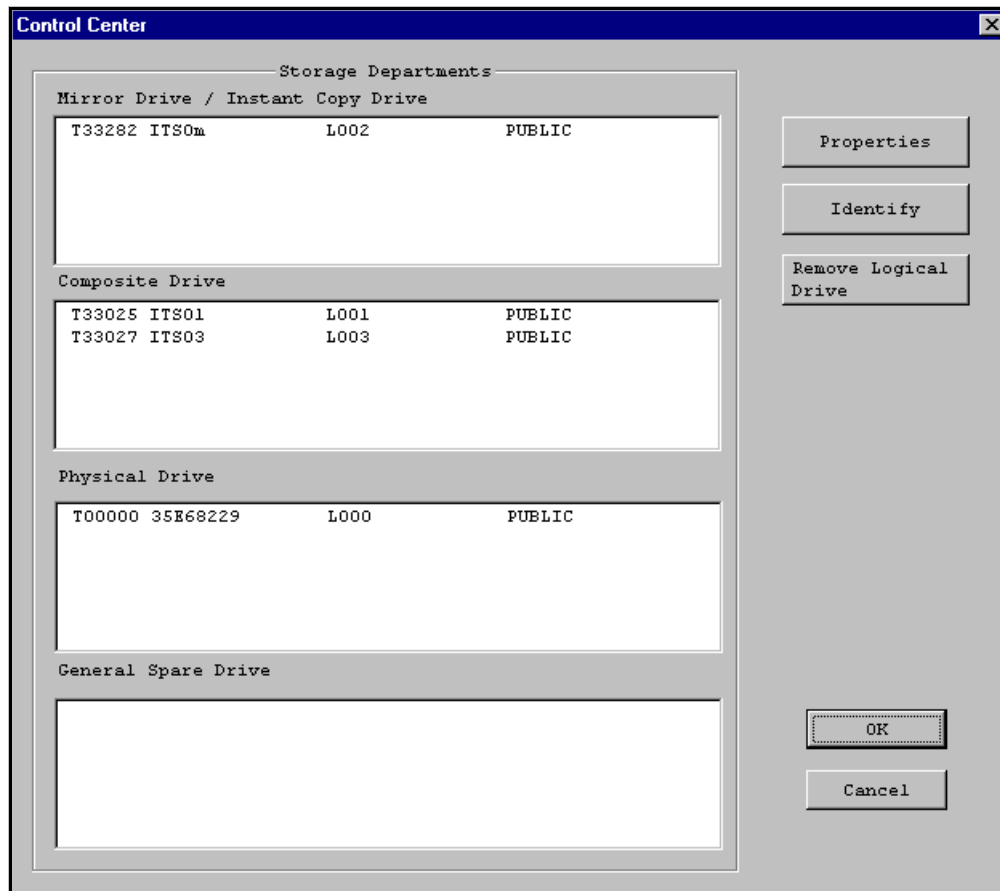


Figure 5-32 Control Center with two composite drives

5.8.2 Creating the mirror

The mirror can now be created by selecting **Tools -> Mirror Drive Setup Wizard**. When the Member Selection window appears, select the composite drives as members of a mirror. Figure 5-33 shows where composite drives 'ITS01' and 'ITS03' are selected as members of a mirror.

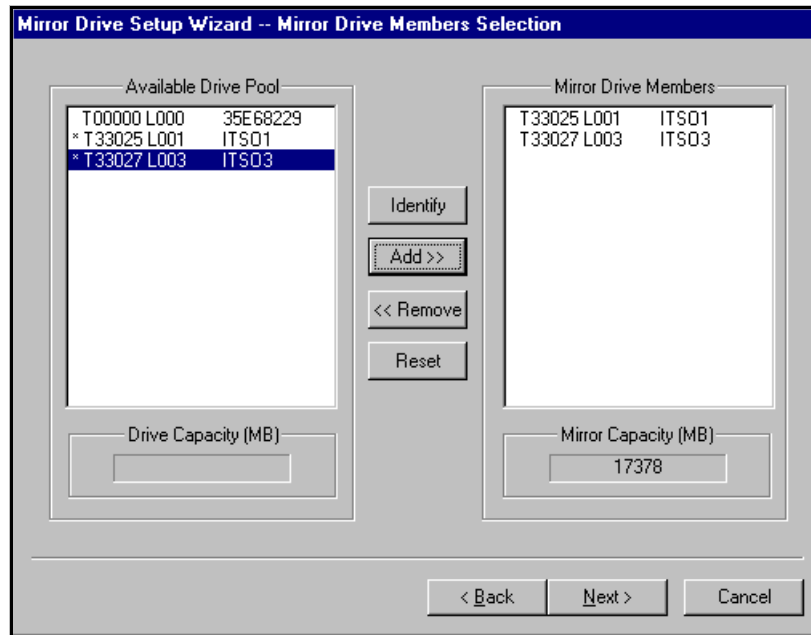


Figure 5-33 Creating mirror drive from two composite drives

Continue through the Mirror Drive Setup Wizard to complete the process as described in 5.6, "Mirror drive" on page 342.

5.8.3 Viewing mirror drive using composite drives

With the Mirror Drive Setup Wizard completed, you can now view the Control Center window once again, as shown in Figure 5-34.

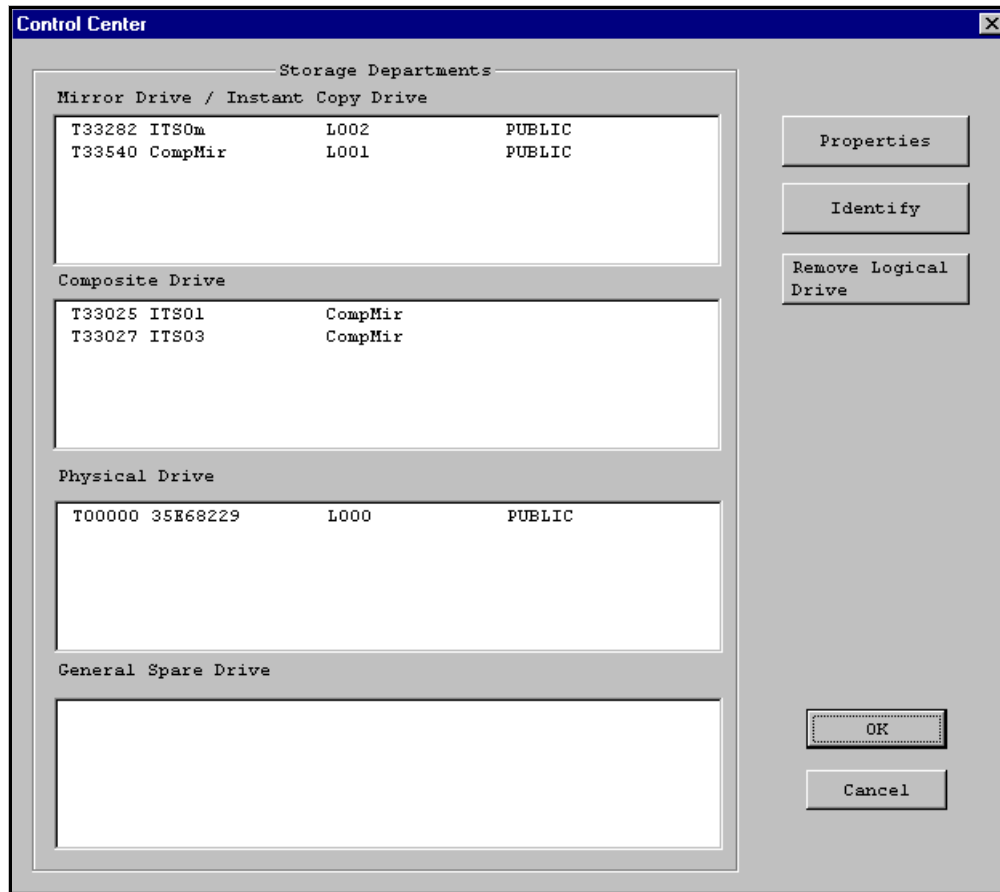


Figure 5-34 Control Center with mirror drive using two composite drives

In the Mirror Drive window, you can see the new mirror drive that was created above and named *CompMir*. In the Composite Drive window you can see that there are still the two composite drives, but instead of having Fibre Channel LUN numbers assigned to them, they are shown as belonging to a mirror with the name *CompMir*.

You can highlight the CompMir drive and click the **Properties** button. All the same functions that were described in 5.6.2, “Mirror drive properties” on page 346 are available.

5.9 Reusing logical drives

At some point the composite, mirror, and instant copy logical drives that have created may be no longer required. The logical drive can be removed so that the member drives that made up the logical drive can then be used individually or reconfigured to make new logical drives.

5.9.1 Remove a logical drive

To remove a logical drive, you access the Control Center by selecting **Tools -> Control Center** from the top toolbar. At the Control Center window, select the logical drive (composite, mirror, or copy) that you want to remove. Select the **Remove Logical Drive** button on the right hand side and a dialog box appears that will ask you to confirm that you want to remove the logical drive.

Once it is removed, the member drives will become general spares and will show up in the General Spare Drive window of the Control Center. This is shown in Figure 5-35.

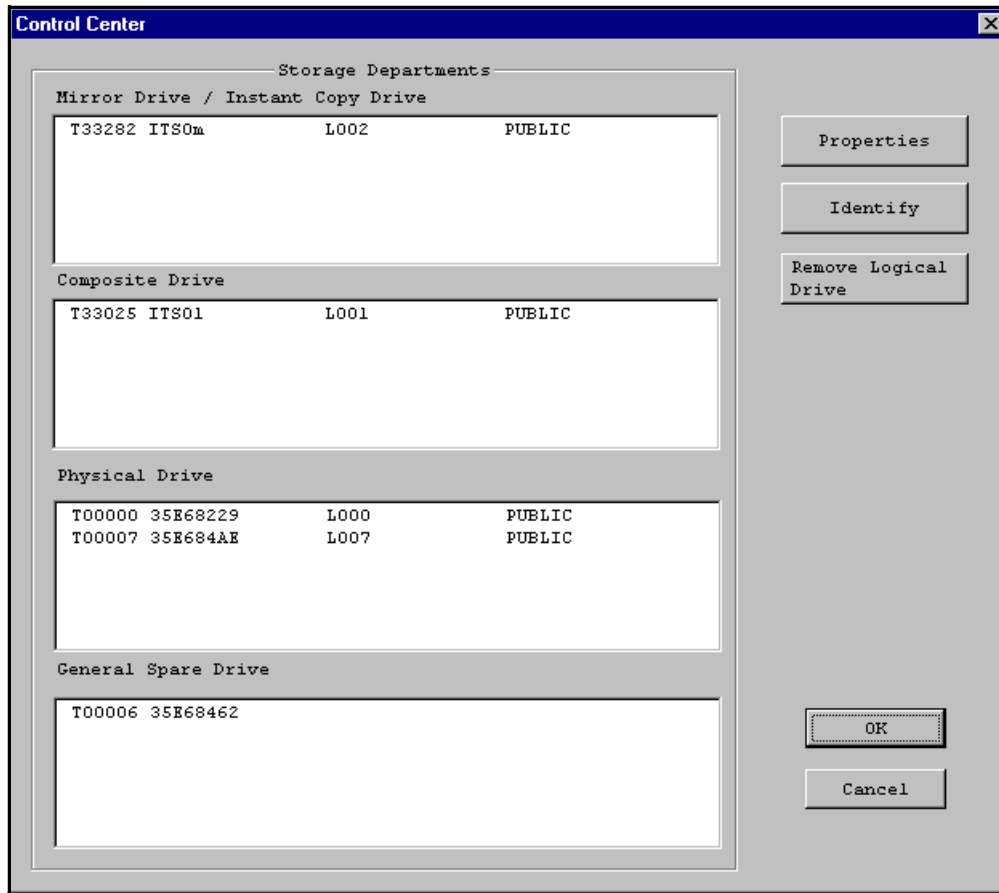


Figure 5-35 Removing a logical drive

The logical drive that was created as an Instant Copy Drive in 5.7.1, “Creating an Instant Copy drive” on page 348, has been removed and is now a general spare.

5.9.2 Mapping a general spare

You will notice in Figure 5-35 that the general spare does not have a LUN number assigned to it. To get a new LUN number for this drive, you select the drive and click the **Properties** button.

The Drive Properties window appears; select the **Change FC** button. A dialog box opens, as shown in Figure 5-36.

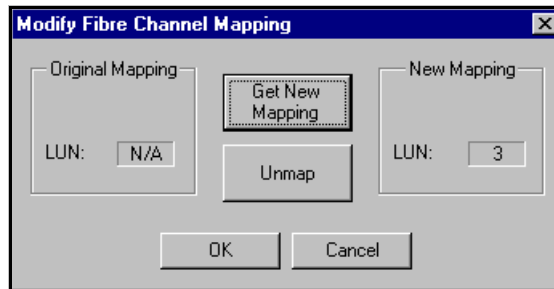


Figure 5-36 Mapping a general spare

The Original Mapping box will indicate that no LUN was assigned previously. Click the **Get New Mapping** button, and the next sequential available LUN number will appear in the New Mapping box. Click **OK**. The drive will now appear in the Physical Drive window of the Control Center.

This can also be used to modify the LUN number on an existing mapped drive, as well as remove the LUN number to 'unmap' a drive and create a general spare.

It is not necessary to map a general spare. A general spare can be used to creating a composite, mirror, or copy drive. Mapping a general spare will create a drive that has a LUN number that can then be used by the host.

5.9.3 Removing a mirror containing composite drive

The mirror in this case was made from logical drives on their own. Once the mirror is removed, the composite drives that made up the mirror will return to the Composite Drive window as viewed from the Control Center.

However, since each composite drive had its attributes changed as it became a member of the mirror, it will no longer be mapped. The composite drives will show up as **UnMapped** in the Control Center window. This is shown in Figure 5-37. The mirror created in 5.8.2, “Creating the mirror” on page 355 was removed.

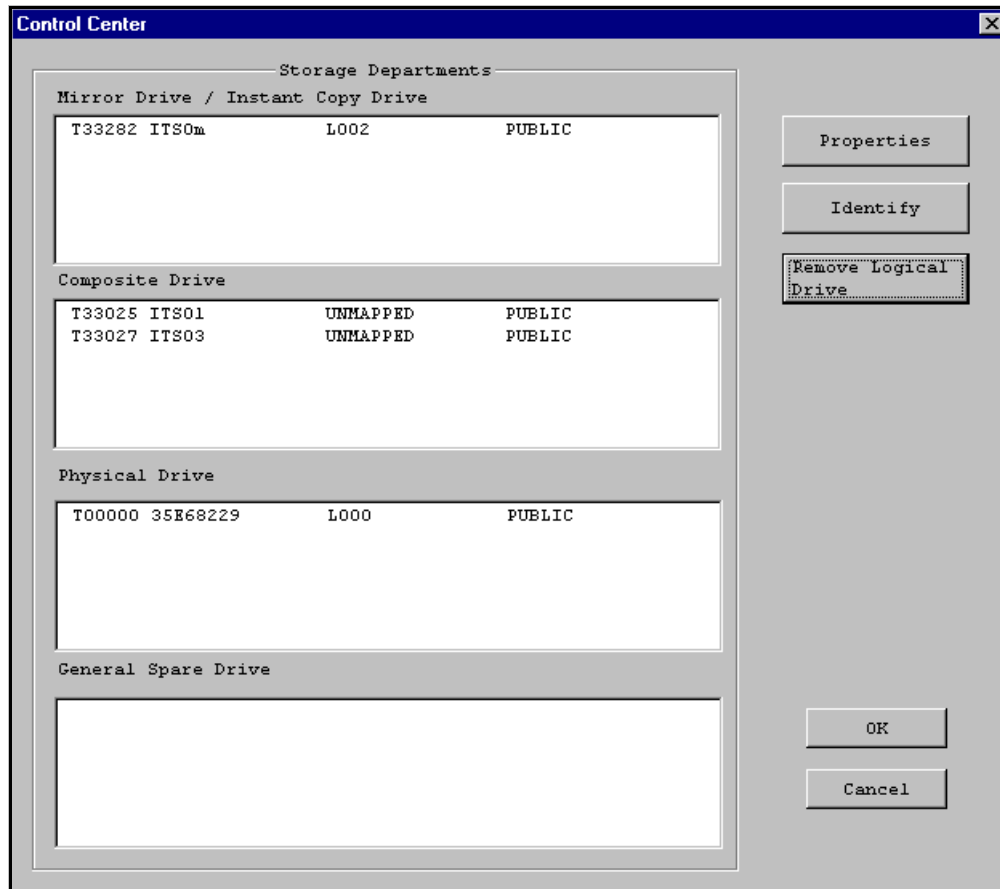


Figure 5-37 UnMapped composite drives

The existing composite drives ‘ITS01’ and ‘ITS03’ cannot be mapped or given a new LUN number at this point. Each logical composite drive *must* be removed as well. This will cause the composite drive to be removed and its member drives to become general spares. Once they are general spares, the drives can be used to recreate the composite drive or to create new logical drives.

5.10 Expanding the SLIC system

The SLIC storage system can be expanded to add more SSA disks or more SLIC Routers. Each SLIC storage system can support up to 64 SSA disks and have 16 Routers.

5.10.1 Adding disk

To add disk to an existing SLIC system is very easy, because they are SSA disks and the rules for SSA disks apply here as well.

If there is a dummy drive in an existing SSA drawer, then it can be replaced by a real drive. If a new drawer has to be added, the SSA cabling is changed to include this new drawer.

Because this is SSA, this can be done *on the fly*. As the SSA loop is broken, the Router will still access all disks due to the structure of the SSA loop. If possible, we recommended that you stop host access and power down the loop. In any case, the rules regarding SSA disks and cabling must be adhered.

As disks are added to an existing loop, the new disks will be recognized. If all disks in the loop are used as single disks (JBOD) and have LUN numbers assigned, the new disks added will have LUN numbers assigned to them automatically. If there are any composite, mirror, instant copy, or spare drives in the loop, then the new disks will not have LUN numbers assigned and become general spares.

5.10.2 Adding Routers

By adding Routers we can increase the amount of storage a host can access and increase throughput. On the rear panel of the Router, there are two Fibre Channel GBIC ports that are available and act as a mini-hub.

You can add a Fibre Channel cable from the second port on the existing Router to one of the ports on the second Router. You are basically daisy-chaining the Routers. But since the ports on the Router act as a hub, an arbitrated loop is created. However, in this scenario there is only one Fibre Channel cable from the host to the Router and it is a single point of failure.

Another option is to add a second Fibre Channel host adapter that will connect to the other Router. This provides a high availability feature, because there are now two paths to the storage system. Software must be used for automatic failover and load balancing between the two Fibre Channel host adapters. Failover also can be done manually if so desired.

On the SSA side, there are a few options available. Each Router can have its own SSA loop so that each one can support 64 SSA disks. In this way, storage capacity is scalable, because it can be increased by adding more Routers. This is shown in Figure 5-38.

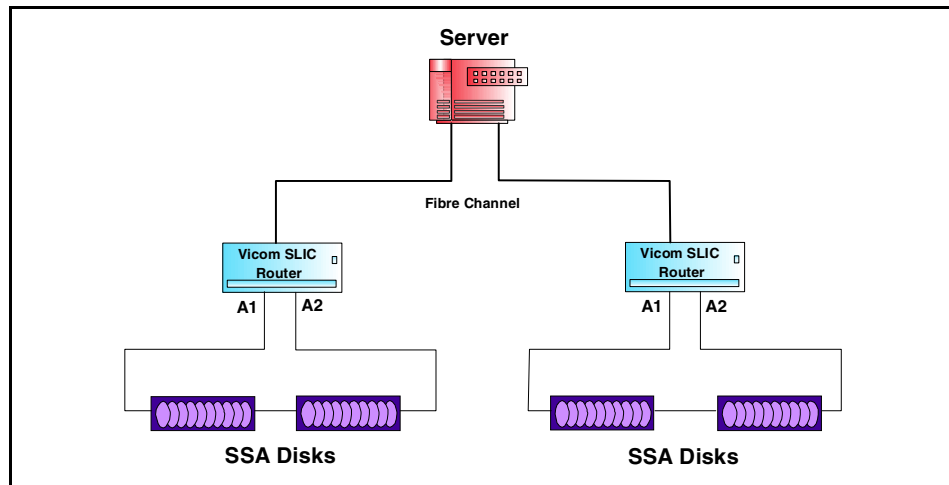


Figure 5-38 Increasing storage capacity

The other option is to have each additional Router added to the same SSA loop. Throughput to the SSA loop will increase, because each Router can access the disks for multiple simultaneous operations. This configuration is shown in Figure 5-39.




Figure 5-39 Increasing throughput

Installing additional Routers

With the existing Router and storage system powered on:

1. The node map on the new Router must be cleared first.
2. Connect the Router to the existing system with the power off.
3. Set SW2 to mode 3 and set SW1 to an unique Fibre Channel target ID.
4. Power on the new Router.
5. When the Status LED on the new Router is on (solid lit), the install is complete.

When the new Router is powered on, communication between the two Routers will occur to query and update the new Router. The new Router will be added as a subordinate so that the first Router will maintain its master status.

Any configuration changes to the storage system is always done on the master. As this is done, the changes are also communicated to the subordinate Routers so that all systems in the loop are aware of what is happening.

Using SLIC Manager on additional Routers

Since the master Router does all the work, it is not necessary to use SLIC Manager to view the new Router. However, depending on how the new Router is used this may become a requirement. The same daemon used to connect to the first Router can be used to connect to several Routers.

The configuration file is edited again and the process to name and create a SLIC zone on this new Router can be added within the same file. The SLIC Manager software can now be used to communicate to the new Router.

The SLIC Manager can only communicate to a single Router at a time. Multiple SLIC Manager sessions can be started to communicate to each Router.

Master failover

If at some time the Router that is designated as the master within the SLIC storage system fails, the master designation will failover to next nearest Router. This is accomplished within the communications between the Routers and it is done automatically.

When the new Router accepts becoming master, it will maintain the master role if even the failed Router is replaced and rejoins the storage system. The master role can be changed back to the original Router, or to another Router if desired, using the SLIC Manager software.

There is no capability to select a specific “failover” Router.

5.10.3 Adding hosts

The SLIC storage system can be expanded to include more hosts whether they are homogeneous or heterogeneous. It is recommended that as hosts are added, each host is connected to its own and separate SLIC Router.

If more than one host was connected to a single Router, there will be arbitration and performance issues. Also, it would have a single point of failure with the possibility of losing data access to many systems.

Homogeneous hosts

If another host of is added and you would like to have both hosts access the same disks, then some sort of access sharing software must be loaded onto both hosts.

If other hosts are added to the storage system and they will not share data, but are connected for storage consolidation, there are a few issues to be considered as the Router does not provide a LUN masking capability.

In UNIX systems, the hosts will see all disk in the loop. But, if the specific volume is not mounted there will be no data integrity problems.

For Windows NT, each host will write its own signature on all available disk. Adding another Windows NT host to the loop will cause problems. To allow a specific Router, and host attached to that Router, access to a specific disk or set of disks, you can set Private Attributes on the disks.

Private Attributes is a setting within SLIC manager that can set a disk to only be accessed by a certain Router and in turn the host attached to that Router.

Note: For more information and operation on the Private Attributes setting, please refer to the SLIC Manager Installation and User Guide, 310-605807

In all cases, if extra control for disk access is required, a third party software, such as Tivoli SANergy, must be used.

Heterogeneous hosts

As the Router does not provide for LUN masking, you must use the SLIC Manager Private Attribute setting or a third party software, such as Tivoli SANergy, to restrict and control host access to the disk. The Private Attributes and Tivoli SANergy can be used together for added control.

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this redbook.

IBM Redbooks

For information on ordering these publications, see “How to get IBM Redbooks” on page 368.

- ▶ *IBM SAN Survival Guide*, SG24-6143
- ▶ *IBM SAN Survival Guide Featuring the INRANGE Portfolio*, SG24-6150
- ▶ *Designing an IBM Storage Area Network*, SG24-5758
- ▶ *Implementing an Open IBM SAN*, SG24-6116
- ▶ *Introduction to Storage Area Network, SAN*, SG24-5470
- ▶ *Introduction to SAN Distance Solutions*, SG24-6408
- ▶ *Introducing Hosts to the SAN Fabric*, SG24-6411
- ▶ *IBM Storage Solutions for Server Consolidation*, SG24-5355
- ▶ *Implementing the Enterprise Storage Server in Your Environment*, SG24-5420
- ▶ *Storage Area Networks: Tape Future In Fabrics*, SG24-5474
- ▶ *IBM Enterprise Storage Server*, SG24-5465

Other resources

These publications are also relevant as further information sources:

- ▶ *Building Storage Networks*, ISBN 0072120509

These IBM publications are also relevant as further information sources:

- ▶ *ESS Web Interface User's Guide for ESS Specialist and ESS Copy Services*, SC26-7346
- ▶ *IBM Storage Area Network Data Gateway Installation and User's Guide*, SC26-7304
- ▶ *IBM Enterprise Storage Server Configuration Planner*, SC26-7353

- ▶ *IBM Enterprise Storage Server Quick Configuration Guide*, SC26-7354
- ▶ *IBM Enterprise Storage Server Introduction and Planning Guide*, 2105 Models E10, E20, F10 and F20, GC26-7294
- ▶ *IBM Enterprise Storage Server User's Guide*, 2105 Models E10, E20, F10 and F20, SC26-7295
- ▶ *IBM Enterprise Storage Server Host Systems Attachment Guide*, 2105 Models E10, E20, F10 and F20, SC26-7296
- ▶ *IBM Enterprise Storage Server SCSI Command Reference*, 2105 Models E10, E20, F10 and F20, SC26-7297
- ▶ *IBM Enterprise Storage Server System/390 Command Reference*, 2105 Models E10, E20, F10 and F20, SC26-7298
- ▶ *IBM Storage Solutions Safety Notices*, GC26-7229
- ▶ *Translated External Devices/Safety Information*, SA26-7003
- ▶ *Electrical Safety for IBM Customer Engineers*, S229-8124
- ▶ *SLIC Router Installation and Users Guide*, 310-605759
- ▶ *SLIC Manager Installation and User Guide*, 310-605807

These INRANGE publications are also relevant as further information sources:

- ▶ *IN-VSN FC/9000 Fibre Channel Director Installation Manual*, 9110461-102
- ▶ *FC/9000 Fibre Channel Director Site Planning Guide*, 9110460-101
- ▶ *FC/9000 Fibre Channel Director Maintenance Manual*, 9110774-307
- ▶ *IN-VSN Enterprise Manager (IN-VSN EM) Software Installation and Operation Guide*, 9110509-203

The JNl publications which are also relevant as further information sources are available on the Web at:

- ▶ <http://www.jni.com/Support/installguides.cfm>

These QLogic publications are also relevant as further information sources:

- ▶ *QLA2200 Hardware Manual*, FC0151103-00
- ▶ *QLA2200 Hardware Manual*, FC0151103-00
- ▶ *QLA2100 Software Manual*, FC0153301-00
- ▶ *QLA2100 Hardware Manual*, FC0151102-00
- ▶ *QMS V1 Installation Guide*, FC0051104-00
- ▶ *QLview for Fibre Operations Guide*, FC0051101-00
- ▶ *QLconfig Operations Guide*, FC0051102-00

Referenced Web sites

These Web sites are also relevant as further information sources:

- ▶ www.storage.ibm.com/ibmsan/index.htm
IBM Enterprise SAN
- ▶ www.pc.ibm.com/ww/netfinity/san
IBM Storage Area Networks: Nefinity Servers
- ▶ www.storage.ibm.com/hardsoft/products/sangateway/supserver.htm
IBM SAN Data Gateway
- ▶ www.storage.ibm.com/hardsoft/products/tape/ro3superserver.htm IBM
SAN Data Gateway Router
- ▶ www.storage.ibm.com/hardsoft/products/fcss/fcss.htm
IBM Fibre Channel RAID Storage Server
- ▶ www.storage.ibm.com/hardsoft/products/ess/ess.htm
Enterprise Storage Server
- ▶ www.cdp.com
Columbia Data Products, Inc.
- ▶ www.emulex.com
Emulex Corporation
- ▶ www.fibrechannel.com
Fibre Channel Industry Association
- ▶ www.jni.com
JNI Corporation
- ▶ www.inrange.com
INRANGE Technologies Corporation
- ▶ www.pathlight.com
Pathlight
- ▶ www.qlogic.com
QLogic Corporation
- ▶ www.sanergy.com
Tivoli SANergy
- ▶ www.snia.org
Storage Networking Industry Association
- ▶ www.tivoli.com
Tivoli
- ▶ www.t11.org
Technical Committee T11

- ▶ www.vicom.com
Vicom Systems
- ▶ www.vixel.com
Vixel
- ▶ www.scsita.org
SCSI Trade Association
- ▶ www.futureio.org
InfiniBand (SM) Trade Association
- ▶ www.nsic.org
National Storage Industry Consortium
- ▶ www.ietf.org
Internet Engineering Task Force
- ▶ www.ansi.org
American National Standards Institute
- ▶ www.standards.ieee.org
Institute of Electrical and Electronics Engineers
- ▶ www.pc.ibm.com/us
US Personal Systems Group

How to get IBM Redbooks

Search for additional Redbooks or redpieces, view, download, or order hardcopy from the Redbooks Web site:

ibm.com/redbooks

Also download additional materials (code samples or diskette/CD-ROM images) from this Redbooks site.

Redpieces are Redbooks in progress; not all Redbooks become redpieces and sometimes just a few chapters will be published this way. The intent is to get the information out much quicker than the formal publishing process allows.

IBM Redbooks collections

Redbooks are also available on CD-ROMs. Click the CD-ROMs button on the Redbooks Web site for information about all the CD-ROMs offered, as well as updates and formats.

Special notices

References in this publication to IBM products, programs or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service may be used. Any functionally equivalent program that does not infringe any of IBM's intellectual property rights may be used instead of the IBM product, program or service.

Information in this book was developed in conjunction with use of the equipment specified, and is limited in application to those specific hardware and software products and levels.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to the IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact IBM Corporation, Dept. 600A, Mail Drop 1329, Somers, NY 10589 USA.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The information contained in this document has not been submitted to any formal IBM test and is distributed AS IS. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

Any pointers in this publication to external Web sites are provided for convenience only and do not in any manner serve as an endorsement of these Web sites.

The following terms are trademarks of other companies:

Tivoli, Manage. Anything. Anywhere., The Power To Manage., Anything. Anywhere., TME, NetView, Cross-Site, Tivoli Ready, Tivoli Certified, Planet Tivoli, and Tivoli Enterprise are trademarks or registered trademarks of Tivoli Systems Inc., an IBM company, in the United States, other countries, or both. In Denmark, Tivoli is a trademark licensed from Kjøbenhavns Sommer - Tivoli A/S.

C-bus is a trademark of Corollary, Inc. in the United States and/or other countries.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and/or other countries.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States and/or other countries.

PC Direct is a trademark of Ziff Communications Company in the United States and/or other countries and is used by IBM Corporation under license.

ActionMedia, LANDesk, MMX, Pentium and ProShare are trademarks of Intel Corporation in the United States and/or other countries.

UNIX is a registered trademark in the United States and other countries licensed exclusively through The Open Group.

SET, SET Secure Electronic Transaction, and the SET Logo are trademarks owned by SET Secure Electronic Transaction LLC.

Other company, product, and service names may be trademarks or service marks of others.

Glossary

8B/10B A data encoding scheme developed by IBM, translating byte-wide data to an encoded 10-bit format. Fibre Channel's FC-1 level defines this as the method to be used to encode and decode data transmissions over the Fibre channel.

Adapter A hardware unit that aggregates other I/O units, devices or communications links to a system bus.

ADSM ADSTAR Distributed Storage Manager.

Agent (1) In the client-server model, the part of the system that performs information preparation and exchange on behalf of a client or server application. (2) In SNMP, the word agent refers to the managed system. See also: Management Agent

AIT Advanced Intelligent Tape - A magnetic tape format by Sony that uses 8mm cassettes, but is only used in specific drives.

AL See Arbitrated Loop

ANSI American National Standards Institute - The primary organization for fostering the development of technology standards in the United States. The ANSI family of Fibre Channel documents provide the standards basis for the Fibre Channel architecture and technology. See FC-PH

Arbitration The process of selecting one respondent from a collection of several candidates that request service concurrently.

Arbitrated Loop A Fibre Channel interconnection technology that allows up to 126 participating node ports and one participating fabric port to communicate.

ATL Automated Tape Library - Large scale tape storage system, which uses multiple tape drives and mechanisms to address 50 or more cassettes.

ATM Asynchronous Transfer Mode - A type of packet switching that transmits fixed-length units of data.

Backup A copy of computer data that is used to recreate data that has been lost, mislaid, corrupted, or erased. The act of creating a copy of computer data that can be used to recreate data that has been lost, mislaid, corrupted or erased.

Bandwidth Measure of the information capacity of a transmission channel.

Bridge (1) A component used to attach more than one I/O unit to a port. (2) A data communications device that connects two or more networks and forwards packets between them. The bridge may use similar or dissimilar media and signaling systems. It operates at the data link level of the OSI model. Bridges read and filter data packets and frames.

Bridge/Router A device that can provide the functions of a bridge, router or both concurrently. A bridge/router can route one or more protocols, such as TCP/IP, and bridge all other traffic. See also: Bridge, Router

Broadcast Sending a transmission to all N_Ports on a fabric.

Channel A point-to-point link, the main task of which is to transport data from one point to another.

Channel I/O A form of I/O where request and response correlation is maintained through some form of source, destination and request identification.

CIFS Common Internet File System

Class of Service A Fibre Channel frame delivery scheme exhibiting a specified set of delivery characteristics and attributes.

Class-1 A class of service providing dedicated connection between two ports with confirmed delivery or notification of non-deliverability.

Class-2 A class of service providing a frame switching service between two ports with confirmed delivery or notification of non-deliverability.

Class-3 A class of service providing frame switching datagram service between two ports or a multicast service between a multicast originator and one or more multicast recipients.

Class-4 A class of service providing a fractional bandwidth virtual circuit between two ports with confirmed delivery or notification of non-deliverability.

Class-6 A class of service providing a multicast connection between a multicast originator and one or more multicast recipients with confirmed delivery or notification of non-deliverability.

Client A software program used to contact and obtain data from a *server* software program on another computer -- often across a great distance. Each *client* program is designed to work specifically with one or more kinds of server programs and each server requires a specific kind of client program.

Client/Server The relationship between machines in a communications network. The client is the requesting machine, the server the supplying machine. Also used to describe the information management relationship between software components in a processing system.

Cluster A type of parallel or distributed system that consists of a collection of interconnected whole computers and is used as a single, unified computing resource.

Coaxial Cable A transmission media (cable) used for high speed transmission. It is called *coaxial* because it includes one physical channel that carries the signal surrounded (after a layer of insulation) by another concentric physical channel, both of which run along the same axis. The inner channel carries the signal and the outer channel serves as a ground.

Controller A component that attaches to the system topology through a channel semantic protocol that includes some form of request/response identification.

CRC Cyclic Redundancy Check - An error-correcting code used in Fibre Channel.

DASD Direct Access Storage Device - any on-line storage device: a disc, drive or CD-ROM.

DAT Digital Audio Tape - A tape media technology designed for very high quality audio recording and data backup. DAT cartridges look like audio cassettes and are often used in mechanical auto-loaders. typically, a DAT cartridge provides 2GB of storage. But new DAT systems have much larger capacities.

Data Sharing A SAN solution in which files on a storage device are shared between multiple hosts.

Datagram Refers to the Class 3 Fibre Channel Service that allows data to be sent rapidly to multiple devices attached to the fabric, with no confirmation of delivery.

dB Decibel - a ratio measurement distinguishing the percentage of signal attenuation between the input and output power. Attenuation (loss) is expressed as dB/km

Disk Mirroring A fault-tolerant technique that writes data simultaneously to two hard disks using the same hard disk controller.

Disk Pooling A SAN solution in which disk storage resources are pooled across multiple hosts rather than be dedicated to a specific host.

DLT Digital Linear Tape - A magnetic tape technology originally developed by Digital Equipment Corporation (DEC) and now sold by Quantum. DLT cartridges provide storage capacities from 10 to 35GB.

E_Port Expansion Port - a port on a switch used to link multiple switches together into a Fibre Channel switch fabric.

ECL Emitter Coupled Logic - The type of transmitter used to drive copper media such as Twinax, Shielded Twisted Pair, or Coax.

Enterprise Network A geographically dispersed network under the auspices of one organization.

Entity In general, a real or existing thing from the Latin *ens*, or being, which makes the distinction between a thing's existence and its qualities. In programming, engineering and probably many other contexts, the word is used to identify units, whether concrete things or abstract ideas, that have no ready name or label.

ESCON Enterprise System Connection

Exchange A group of sequences which share a unique identifier. All sequences within a given exchange use the same protocol. Frames from multiple sequences can be multiplexed to prevent a single exchange from consuming all the bandwidth. See also: Sequence

F_Node Fabric Node - a fabric attached node.

F_Port Fabric Port - a port used to attach a Node Port (N_Port) to a switch fabric.

Fabric Fibre Channel employs a fabric to connect devices. A fabric can be as simple as a single cable connecting two devices. The term is most often used to describe a more complex network utilizing hubs, switches and gateways.

Fabric Login Fabric Login (FLOGI) is used by an N_Port to determine if a fabric is present and, if so, to initiate a session with the fabric by exchanging service parameters with the fabric. Fabric Login is performed by an N_Port following link initialization and before communication with other N_Ports is attempted.

FC Fibre Channel

FC-0 Lowest level of the Fibre Channel Physical standard, covering the physical characteristics of the interface and media

FC-1 Middle level of the Fibre Channel Physical standard, defining the 8B/10B encoding/decoding and transmission protocol.

FC-2 Highest level of the Fibre Channel Physical standard, defining the rules for signaling protocol and describing transfer of frame, sequence and exchanges.

FC-3 The hierarchical level in the Fibre Channel standard that provides common services such as striping definition.

FC-4 The hierarchical level in the Fibre Channel standard that specifies the mapping of upper-layer protocols to levels below.

FCA Fiber Channel Association.

FC-AL Fibre Channel Arbitrated Loop - A reference to the Fibre Channel Arbitrated Loop standard, a shared gigabit media for up to 127 nodes, one of which may be attached to a switch fabric. See also: Arbitrated Loop.

FC-CT Fibre Channel common transport protocol

FC-FG Fibre Channel Fabric Generic - A reference to the document (ANSI X3.289-1996) which defines the concepts, behavior and characteristics of the Fibre Channel Fabric along with suggested partitioning of the 24-bit address space to facilitate the routing of frames.

FC-FP Fibre Channel HIPPI Framing Protocol - A reference to the document (ANSI X3.254-1994) defining how the HIPPI framing protocol is transported via the fibre channel

FC-GS Fibre Channel Generic Services - A reference to the document (ANSI X3.289-1996) describing a common transport protocol used to communicate with the server functions, a full X500 based directory service, mapping of the Simple Network Management Protocol (SNMP) directly to the Fibre Channel, a time server and an alias server.

FC-LE Fibre Channel Link Encapsulation - A reference to the document (ANSI X3.287-1996) which defines how IEEE 802.2 Logical Link Control (LLC) information is transported via the Fibre Channel.

FC-PH A reference to the Fibre Channel Physical and Signaling standard ANSI X3.230, containing the definition of the three lower levels (FC-0, FC-1, and FC-2) of the Fibre Channel.

FC-PLDA Fibre Channel Private Loop Direct Attach - See PLDA.

FC-SB Fibre Channel Single Byte Command Code Set - A reference to the document (ANSI

X.271-1996) which defines how the ESCON command set protocol is transported using the fibre channel.

FC-SW Fibre Channel Switch Fabric - A reference to the ANSI standard under development that further defines the fabric behavior described in FC-FG and defines the communications between different fabric elements required for those elements to coordinate their operations and management address assignment.

FC Storage Director See SAN Storage Director

FCA Fibre Channel Association - a Fibre Channel industry association that works to promote awareness and understanding of the Fibre Channel technology and its application and provides a means for implementers to support the standards committee activities.

FCLC Fibre Channel Loop Association - an independent working group of the Fibre Channel Association focused on the marketing aspects of the Fibre Channel Loop technology.

FCP Fibre Channel Protocol - the mapping of SCSI-3 operations to Fibre Channel.

Fiber Optic Refers to the medium and the technology associated with the transmission of information along a glass or plastic wire or fiber.

Fibre Channel A technology for transmitting data between computer devices at a data rate of up to 4 Gb/s. It is especially suited for connecting computer servers to shared storage devices and for interconnecting storage controllers and drives.

FICON Fibre Connection - A next-generation I/O solution for IBM S/390 parallel enterprise server.

FL_Port Fabric Loop Port - the access point of the fabric for physically connecting the user's Node Loop Port (NL_Port).

FLOGI See Fabric Log In

Frame A linear set of transmitted bits that define the basic transport unit. The frame is the most basic element of a message in Fibre Channel communications, consisting of a 24-byte header and zero to 2112 bytes of data. See also: Sequence

FSP Fibre Channel Service Protocol - The common FC-4 level protocol for all services, transparent to the fabric type or topology.

Full-Duplex A mode of communications allowing simultaneous transmission and reception of frames.

G_Port Generic Port - a generic switch port that is either a Fabric Port (F_Port) or an Expansion Port (E_Port). The function is automatically determined during login.

Gateway A node on a network that interconnects two otherwise incompatible networks.

Gb/s Gigabits per second. Also sometimes referred to as Gbps. In computing terms it is approximately 1,000,000,000 bits per second. Most precisely it is 1,073,741,824 (1024 x 1024 x 1024) bits per second.

GB/s Gigabytes per second. Also sometimes referred to as GBps. In computing terms it is approximately 1,000,000,000 bytes per second. Most precisely it is 1,073,741,824 (1024 x 1024 x 1024) bytes per second.

GBIC GigaBit Interface Converter - Industry standard transceivers for connection of Fibre Channel nodes to arbitrated loop hubs and fabric switches.

Gigabit One billion bits, or one thousand megabits.

GLM Gigabit Link Module - a generic Fibre Channel transceiver unit that integrates the key functions necessary for installation of a Fibre channel media interface on most systems.

Half-Duplex A mode of communications allowing either transmission or reception of frames at any point in time, but not both (other than link control frames which are always permitted).

Hardware The mechanical, magnetic and electronic components of a system, e.g., computers, telephone switches, terminals and the like.

HBA Host Bus Adapter

HIPPI High Performance Parallel Interface - An ANSI standard defining a channel that transfers

data between CPUs and from a CPU to disk arrays and other peripherals.

HMMP HyperMedia Management Protocol

HMMS HyperMedia Management Schema - the definition of an implementation-independent, extensible, common data description/schema allowing data from a variety of sources to be described and accessed in real time regardless of the source of the data. See also: WEBM, HMMP

HSM Hierarchical Storage Management - A software and hardware system that moves files from disk to slower, less expensive storage media based on rules and observation of file activity. Modern HSM systems move files from magnetic disk to optical disk to magnetic tape.

HUB A Fibre Channel device that connects nodes into a logical loop by using a physical star topology. Hubs will automatically recognize an active node and insert the node into the loop. A node that fails or is powered off is automatically removed from the loop.

HUB Topology see Loop Topology

Hunt Group A set of associated Node Ports (N_Ports) attached to a single node, assigned a special identifier that allows any frames containing this identifier to be routed to any available Node Port (N_Port) in the set.

In-Band Signaling This is signaling that is carried in the same channel as the information. Also referred to as inband.

Information Unit A unit of information defined by an FC-4 mapping. Information Units are transferred as a Fibre Channel Sequence.

Intermix A mode of service defined by Fibre Channel that reserves the full Fibre Channel bandwidth for a dedicated Class 1 connection, but also allows connection-less Class 2 traffic to share the link if the bandwidth is available.

I/O Input/output

IP Internet Protocol

IPI Intelligent Peripheral Interface

Isochronous Transmission Data transmission which supports network-wide timing

requirements. A typical application for isochronous transmission is a broadcast environment which needs information to be delivered at a predictable time.

JBOD Just a bunch of disks.

Jukebox A device that holds multiple optical disks and one or more disk drives, and can swap disks in and out of the drive as needed.

L_Port Loop Port - A node or fabric port capable of performing Arbitrated Loop functions and protocols. NL_Ports and FL_Ports are loop-capable ports.

LAN See Local Area Network - A network covering a relatively small geographic area (usually not larger than a floor or small building). Transmissions within a Local Area Network are mostly digital, carrying data among stations at rates usually above one megabit/s.

Latency A measurement of the time it takes to send a frame between two locations.

LC Lucent Connector. A registered Trademark of Lucent Technologies.

Link A connection between two Fibre Channel ports consisting of a transmit fibre and a receive fibre.

Link_Control_Facility A termination card that handles the logical and physical control of the Fibre Channel link for each mode of use.

LIP A Loop Initialization Primitive sequence is a special fibre channel sequence that is used to start loop initialization. Allows ports to establish their port addresses.

Local Area Network (LAN) A network covering a relatively small geographic area (usually not larger than a floor or small building). Transmissions within a Local Area Network are mostly digital, carrying data among stations at rates usually above one megabit/s.

Login Server Entity within the Fibre Channel fabric that receives and responds to login requests.

Loop Circuit A temporary point-to-point like path that allows bi-directional communications between loop-capable ports.

Loop Topology An interconnection structure in which each point has physical links to two neighbors resulting in a closed circuit. In a loop topology, the available bandwidth is shared.

LVD Low Voltage Differential

Management Agent A process that exchanges a managed node's information with a management station.

Managed Node A managed node is a computer, a storage system, a gateway, a media device such as a switch or hub, a control instrument, a software product such as an operating system or an accounting package, or a machine on a factory floor, such as a robot.

Managed Object A variable of a managed node. This variable contains one piece of information about the node. Each node can have several objects.

Management Station A host system that runs the management software.

Mb/s Megabits per second. Also sometimes referred to as Mbps. In computing terms it is approximately 1,000,000 bits per second. Most precisely it is 1,048,576 (1024 x 1024) bits per second.

MB/s Megabytes per second. Also sometimes referred to as MBps. In computing terms it is approximately 1,000,000 bytes per second. Most precisely it is 1,048,576 (1024 x 1024) bytes per second.

Meter 39.37 inches, or just slightly larger than a yard (36 inches)

Media Plural of medium. The physical environment through which transmission signals pass. Common media include copper and fiber optic cable.

Media Access Rules (MAR).

MIA Media Interface Adapter - MIAs enable optic-based adapters to interface to

copper-based devices, including adapters, hubs, and switches.

MIB Management Information Block - A formal description of a set of network objects that can be managed using the Simple Network Management Protocol (SNMP). The format of the MIB is defined as part of SNMP and is a hierarchical structure of information relevant to a specific device, defined in object oriented terminology as a collection of objects, relations, and operations among objects.

Mirroring The process of writing data to two separate physical devices simultaneously.

MM Multi-Mode - See Multi-Mode Fiber

MMF See Multi-Mode Fiber - - In optical fiber technology, an optical fiber that is designed to carry multiple light rays or modes concurrently, each at a slightly different reflection angle within the optical core. Multi-Mode fiber transmission is used for relatively short distances because the modes tend to disperse over longer distances. See also: Single-Mode Fiber, SMF

Multicast Sending a copy of the same transmission from a single source device to multiple destination devices on a fabric. This includes sending to all N_Ports on a fabric (broadcast) or to only a subset of the N_Ports on a fabric (multicast).

Multi-Mode Fiber (MMF) In optical fiber technology, an optical fiber that is designed to carry multiple light rays or modes concurrently, each at a slightly different reflection angle within the optical core. Multi-Mode fiber transmission is used for relatively short distances because the modes tend to disperse over longer distances. See also: Single-Mode Fiber

Multiplex The ability to intersperse data from multiple sources and destinations onto a single transmission medium. Refers to delivering a single transmission to multiple destination Node Ports (N_Ports).

N_Port Node Port - A Fibre Channel-defined hardware entity at the end of a link which provides the mechanisms necessary to transport information units to or from another node.

N_Port Login N_Port Login (PLOGI) allows two N_Ports to establish a session and exchange identities and service parameters. It is performed following completion of the fabric login process and prior to the FC-4 level operations with the destination port. N_Port Login may be either explicit or implicit.

Name Server Provides translation from a given node name to one or more associated N_Port identifiers.

NAS Network Attached Storage - a term used to describe a technology where an integrated storage system is attached to a messaging network that uses common communications protocols, such as TCP/IP.

NDMP Network Data Management Protocol

Network An aggregation of interconnected nodes, workstations, file servers, and/or peripherals, with its own protocol that supports interaction.

Network Topology Physical arrangement of nodes and interconnecting communications links in networks based on application requirements and geographical distribution of users.

NFS Network File System - A distributed file system in UNIX developed by Sun Microsystems which allows a set of computers to cooperatively access each other's files in a transparent manner.

NL_Port Node Loop Port - a node port that supports Arbitrated Loop devices.

NMS Network Management System - A system responsible for managing at least part of a network. NMSs communicate with agents to help keep track of network statistics and resources.

Node An entity with one or more N_Ports or NL_Ports.

Non-Blocking A term used to indicate that the capabilities of a switch are such that the total number of available transmission paths is equal to the number of ports. Therefore, all ports can have simultaneous access through the switch.

Non-L_Port A Node or Fabric port that is not capable of performing the Arbitrated Loop

functions and protocols. N_Ports and F_Ports are not loop-capable ports.

Operation A term defined in FC-2 that refers to one of the Fibre Channel *building blocks* composed of one or more, possibly concurrent, exchanges.

Optical Disk A storage device that is written and read by laser light.

Optical Fiber A medium and the technology associated with the transmission of information as light pulses along a glass or plastic wire or fiber.

Ordered Set A Fibre Channel term referring to four 10-bit characters (a combination of data and special characters) providing low-level link functions, such as frame demarcation and signaling between two ends of a link.

Originator A Fibre Channel term referring to the initiating device.

Out of Band Signaling This is signaling that is separated from the channel carrying the information. Also referred to as outband.

Peripheral Any computer device that is not part of the essential computer (the processor, memory and data paths) but is situated relatively close by. A near synonym is input/output (I/O) device.

Petard A device that is small and sometimes explosive.

PLDA Private Loop Direct Attach - A technical report which defines a subset of the relevant standards suitable for the operation of peripheral devices such as disks and tapes on a private loop.

PLOGI See N_Port Login

Point-to-Point Topology An interconnection structure in which each point has physical links to only one neighbor resulting in a closed circuit. In point-to-point topology, the available bandwidth is dedicated.

Port The hardware entity within a node that performs data communications over the Fibre Channel.

Port Bypass Circuit A circuit used in hubs and disk enclosures to automatically open or close the loop to add or remove nodes on the loop.

Private NL_Port An NL_Port which does not attempt login with the fabric and only communicates with other NL Ports on the same loop.

Protocol A data transmission convention encompassing timing, control, formatting and data representation.

Public NL_Port An NL_Port that attempts login with the fabric and can observe the rules of either public or private loop behavior. A public NL_Port may communicate with both private and public NL_Ports.

Quality of Service (QoS) A set of communications characteristics required by an application. Each QoS defines a specific transmission priority, level of route reliability, and security level.

RAID Redundant Array of Inexpensive or Independent Disks. A method of configuring multiple disk drives in a storage subsystem for high availability and high performance.

Raid 0 Level 0 RAID support - Striping, no redundancy

Raid 1 Level 1 RAID support - mirroring, complete redundancy

Raid 5 Level 5 RAID support, Striping with parity

Repeater A device that receives a signal on an electromagnetic or optical transmission medium, amplifies the signal, and then retransmits it along the next leg of the medium.

Responder A Fibre Channel term referring to the answering device.

Router (1) A device that can decide which of several paths network traffic will follow based on some optimal metric. Routers forward packets from one network to another based on network-layer information. (2) A dedicated computer hardware and/or software package which manages the connection between two or more networks. See also: Bridge, Bridge/Router

SAF-TE SCSI Accessed Fault-Tolerant Enclosures

SAN A Storage Area Network (SAN) is a dedicated, centrally managed, secure information infrastructure, which enables any-to-any interconnection of servers and storage systems.

SAN System Area Network - term originally used to describe a particular symmetric multiprocessing (SMP) architecture in which a switched interconnect is used in place of a shared bus. Server Area Network - refers to a switched interconnect between multiple SMPs.

SC Connector Subscriber Connector. A fiber optic connector standardized by ANSI TIA/EIA-568A for use in structured wiring installations.

Scalability The ability of a computer application or product (hardware or software) to continue to function well as it (or its context) is changed in size or volume. For example, the ability to retain performance levels when adding additional processors, memory and/or storage.

SCSI Small Computer System Interface - A set of evolving ANSI standard electronic interfaces that allow personal computers to communicate with peripheral hardware such as disk drives, tape drives, CD_ROM drives, printers and scanners faster and more flexibly than previous interfaces. The table below identifies the major characteristics of the different SCSI version.

SCSI Ver- sion	Sig- nal Rate MHz	Bus- Width (bits)	Max. DTR (MB/s)	Max. Num. Devic es	Max. Cable Lengt h (m)
SCSI -1	5	8	5	7	6
SCSI -2	5	8	5	7	6
Wide SCSI -2	5	16	10	15	6
Fast SCSI -2	10	8	10	7	6

Fast Wide SCSI -2	10	16	20	15	6
Ultra SCSI	20	8	20	7	1.5
Ultra SCSI -2	20	16	40	7	12
Ultra 2 LVD SCSI	40	16	80	15	12

SCSI-3 SCSI-3 consists of a set of primary commands and additional specialized command sets to meet the needs of specific device types. The SCSI-3 command sets are used not only for the SCSI-3 parallel interface but for additional parallel and serial protocols, including Fibre Channel, Serial Bus Protocol (used with IEEE 1394 Firewire physical protocol) and the Serial Storage Protocol (SSP).

SCSI-FCP The term used to refer to the ANSI Fibre Channel Protocol for SCSI document (X3.269-199x) that describes the FC-4 protocol mappings and the definition of how the SCSI protocol and command set are transported using a Fibre Channel interface.

Sequence A series of frames strung together in numbered order which can be transmitted over a Fibre Channel connection as a single operation. See also: Exchange

SERDES Serializer Deserializer

Server A computer which is dedicated to one task.

SES SCSI Enclosure Services - ANSI SCSI-3 proposal that defines a command set for soliciting basic device status (temperature, fan speed, power supply status, etc.) from a storage enclosures.

Single-Mode Fiber In optical fiber technology, an optical fiber that is designed for the transmission of a single ray or mode of light as a carrier. It is a single light path used for long-distance signal transmission. See also: Multi-Mode Fiber

SMART Self Monitoring and Reporting Technology

SM Single Mode - See Single-Mode Fiber

SMF Single-Mode Fiber - In optical fiber technology, an optical fiber that is designed for the transmission of a single ray or mode of light as a carrier. It is a single light path used for long-distance signal transmission. See also: MMF

SNIA Storage Networking Industry Association. A non-profit organization comprised of more than 77 companies and individuals in the storage industry.

SN Storage Network. See also: SAN

SNMP Simple Network Management Protocol - The Internet network management protocol which provides a means to monitor and set network configuration and run-time parameters.

SNMWG Storage Network Management Working Group is chartered to identify, define and support open standards needed to address the increased management requirements imposed by storage area network environments.

SSA Serial Storage Architecture - A high speed serial loop-based interface developed as a high speed point-to-point connection for peripherals, particularly high speed storage arrays, RAID and CD-ROM storage by IBM.

ST Straight Tip connector. Also known as "Stick and Twist".

Star The physical configuration used with hubs in which each user is connected by communications links radiating out of a central hub that handles all communications.

StorWatch Expert These are StorWatch applications that employ a 3 tiered architecture that includes a management interface, a StorWatch manager and agents that run on the storage resource(s) being managed. Expert products employ a StorWatch data base that can be used for saving key management data (e.g. capacity or performance metrics). Expert products use the agents as well as analysis of storage data saved in the data base to perform

higher value functions including -- reporting of capacity, performance, etc. over time (trends), configuration of multiple devices based on policies, monitoring of capacity and performance, automated responses to events or conditions, and storage related data mining.

StorWatch Specialist A StorWatch interface for managing an individual fibre Channel device or a limited number of like devices (that can be viewed as a single group). StorWatch specialists typically provide simple, point-in-time management functions such as configuration, reporting on asset and status information, simple device and event monitoring, and perhaps some service utilities.

Striping A method for achieving higher bandwidth using multiple N_Ports in parallel to transmit a single information unit across multiple levels.

STP Shielded Twisted Pair

Storage Media The physical device itself, onto which data is recorded. Magnetic tape, optical disks, floppy disks are all storage media.

Switch A component with multiple entry/exit points (ports) that provides dynamic connection between any two of these points.

Switch Topology An interconnection structure in which any entry point can be dynamically connected to any exit point. In a switch topology, the available bandwidth is scalable.

T11 A technical committee of the National Committee for Information Technology Standards, titled T11 I/O Interfaces. It is tasked with developing standards for moving data in and out of computers.

Tape Backup Making magnetic tape copies of hard disk and optical disc files for disaster recovery.

Tape Pooling A SAN solution in which tape resources are pooled and shared across multiple hosts rather than being dedicated to a specific host.

TCP Transmission Control Protocol - a reliable, full duplex, connection-oriented end-to-end transport protocol running on top of IP.

TCP/IP Transmission Control Protocol/ Internet Protocol - a set of communications protocols that support peer-to-peer connectivity functions for both local and wide area networks.

Time Server A Fibre Channel-defined service function that allows for the management of all timers used within a Fibre Channel system.

Topology An interconnection scheme that allows multiple Fibre Channel ports to communicate. For example, point-to-point, Arbitrated Loop, and switched fabric are all Fibre Channel topologies.

T_Port An ISL port more commonly known as an E_Port, referred to as a Trunk port and used by INRANGE.

TL_Port A private to public bridging of switches or directors, referred to as Translative Loop.

Twinax A transmission media (cable) consisting of two insulated central conducting leads of coaxial cable.

Twisted Pair A transmission media (cable) consisting of two insulated copper wires twisted around each other to reduce the induction (thus interference) from one wire to another. The twists, or lays, are varied in length to reduce the potential for signal interference between pairs. Several sets of twisted pair wires may be enclosed in a single cable. This is the most common type of transmission media.

ULP Upper Level Protocols

UTC Under-The-Covers, a term used to characterize a subsystem in which a small number of hard drives are mounted inside a higher function unit. The power and cooling are obtained from the system unit. Connection is by parallel copper ribbon cable or pluggable backplane, using IDE or SCSI protocols.

UTP Unshielded Twisted Pair

Virtual Circuit A unidirectional path between two communicating N_Ports that permits fractional bandwidth.

WAN Wide Area Network - A network which encompasses inter-connectivity between devices over a wide geographic area. A wide area network may be privately owned or rented, but the term usually connotes the inclusion of public (shared) networks.

WDM Wave Division Multiplexing - A technology that puts data from different sources together on an optical fiber, with each signal carried on its own separate light wavelength. Using WDM, up to 80 (and theoretically more) separate wavelengths or channels of data can be multiplexed into a stream of light transmitted on a single optical fiber.

WEBM Web-Based Enterprise Management - A consortium working on the development of a series of standards to enable active management and monitoring of network-based elements.

Zoning In Fibre Channel environments, the grouping together of multiple ports to form a virtual private storage network. Ports that are members of a group or zone can communicate with each other but are isolated from ports in other zones.

Index

Numerics

1-way mirror 342
2042-001 270
2766 adapter 10
2-way mirror 342
3527 321
3-way mirror 342
7131 321
7133 321

A

access levels 94
access profile 163
Access_Any 163
Access_Restricted 163
activate hard zone layout 252
adapter types 2
Add Volumes 126
Adding an Instant Copy Drive to a Mirror 350
Adding Disk 361
Adding Hosts 364
Adding Routers 361
Adding SCSI devices 295
administrator 279
Advanced Options 286
AdvFS 22
Agent 277
AIX host system 15
All or Nothing 240
anonymous host 163
Arbitrated Loop 228
arbitrated loop 167
Assign Fibre Channel Target 324
Assigned Volumes 104
Assigning LUN IDs 295
Attaching hosts 313
Auto Learning 233
Auto Sense 222
autonomy of routing 270

B

bandwidth 143

bridge 16
bridging 380
broadcast messages 241
bypassed device 227

C

call home 86
cascaded directors 264
channel
 access mode 162
channel command 5
channel path group 6
Client 277
clustering 25
Combining Composite And Mirroring 352
Compaq Alpha server 18
Compaq host FC attachment 18
Compaq Tru64 UNIX 19
composite drive 337, 352
 as members of a mirror 355
 properties 341
Configure Host Adapter Ports 159
Control Center 333
Copy Services 90
Create General Spares 325
create user for telnet to San Data Gateway 275
Creating A Composite Drive 337
Creating A Mirror Drive 342
Creating a SLIC Zone 330
Creating an Instant Copy drive 348

D

daisy chained 270
Detach Instant Copy Drive from a Mirror 352
Device discovery 294
device ports 270
director clock 218
director FIOs 270
disk group 113
disk groups 115, 132
disk pooling 75
distances 3

E

- Emulex 46
- Emulex LP8000 32, 42, 47, 53, 56, 58, 65, 67
- enabling the Ethernet port 275
- Enabling VPS 297
- ESCON 5, 104
- ESS 317–318
- ESS configuration update 110
- ESS configuring 89
- ESS Copy Server 90
- ESS default user 81
- ESS disk group section 100
- ESS FC port 98
- ESS FC port access 123
- ESS FC port definitions 116
- ESS FC ports config 158
- ESS FC ports topology 159
- ESS host interfaces 97
- ESS interface adapters 160
- ESS introduction panel 78
- ESS logical volume 130
- ESS logical volume assignments 142
- ESS logical volumes 126
- ESS Open System Storage 104
- ESS operating 82
- ESS ports 99
- ESS site certificate 79
- ESS status 83
- ESS storage allocation 103
- ESS storage configuration 96
- ESS Storage Server Attributes 162
- ESS to FICON 7
- ESS user administration 93
- ESS Web Copy Services 91
- ESS Web Interface 77

F

- fabric add 211
- FC adapters 98
- FC/9000 7
- FC/9000 overview 177
- FC_AL 237
- FC-AL 167, 220
- FCM blade 267
- FC-PP 167
- FC-SW 167
- feature code 2314 7
- feature code 2315 7

- feature code 2316 7
- feature code 2318 7
- feature Code 3019 16
- feature code 3021 3
- feature code 3022 2
- feature code 3023 3
- Fibre Channel 287
- Fibre Channel adapters 98
- Fibre Channel host 285
- FICON 4–5, 7
- FICON switch cascading 7
- FIO blade 238
- firmware 183
- Fixed Block 109
- Fixed Block Storage 108, 114
- FlashCopy 77
- Format the Drives 325
- frame multiplexing 6
- full-duplex 1

G

- Get New Mapping 359

H

- hard zone setup 251
- Hard zoning 236
- hard zoning 237
- hard zoning setup 247
- HBA configuration files 41
- Heterogeneous Hosts 364
- Hewlett Packard 9000 23
- Hewlett Packard 9000 FC attachment 23
- High availability considerations 317
- Homogeneous Hosts 364
- hop 265
- host adapter 1
- Host Attributes 118
- Host information 302
- Host registration 298
- Hosts Systems 104
- HOSTSW 298

I

- I/O 6
- IBM Enterprise Storage Server, 2105-F20 75
- IBM SAN Fibre Channel Switch 307, 316
- IBM Storage Area Network Data Gateway,

- 2108-G07 271
- Initialization Method 339
- initializeBox 274
- INRANGE defining hard zoning 246
- INRANGE FC/9000 175
- INRANGE hard zoning 237
- INRANGE hard zoning rules 238
- INRANGE Name Server zoning 242
- INRANGE zoning methods 237
- Installing additional Routers 363
- Installing StorWatch Specialist 278
- Installing the SLIC Manager 327
- Installing the SLIC Router 323
- Instant Copy 322, 347
- Instant Copy Drive 342
- Instant Copy Drive Properties 349
- Inter Switch Link 263
- interconnected directors 7
- IN-VSN event log 269
- IOC-0210-54 adapter 17
- IOC-0210-54 adapter card 15, 17
- IP settings 267
- iSeries configurations 10
- iSeries FC host 8
- iSeries host system 8
- ISL 263, 380
- ITSO environment
 - installation steps 176

J

- JBOD 109, 115
- JNI adapters 45
- JNI FC64-1063 42
- JNI FCI-1063 42
- JNI PCI adapter 35
- JNI SBUS 36
- Just a Bunch Of Disks 109

L

- legacy FC 220
- LIC 88
- LIC level 164
- Licensed Internal Code 88
- link extender 3
- LIP 236
- locating WWPN 69
- logical drives 322, 357
- logical subsystems 4

- logical unit number 162
- login 81
- long-wave ports 273
- Loop 308
- loop attachments 221
- loop devices 226
- Loop Initialization Primitive 236
- loop ports 225
- Loop protocol 220
- LP8000 32
- LUN 162
- LUN access 315
- LUN affinity 4
- LUN masking 254, 273, 305, 319, 364
- LUN support 294
- LUNs 4

M

- Mapping a general spare 359
- Mapping Physical Drives 324
- mapRebuildDatabase 295
- mapShowDatabase 304
- mapWinnowDatabase 295
- Master Failover 363
- Master Router 335
- maximum distance 3
- MC/ServiceGuard 25
- Mirror Capacity 344
- mirror capacity 352
- Mirror drive 352
- mirror drive 342
- Mirror Drive Dedicated Spare 344
- Mirror Drive Properties 346
- Modify Host Systems 117
- Modify Users 94
- monitoring an INRANGE SAN 266
- multipathing 256
- multiple controllers 6
- multiple Director FRUs 270

N

- name server table 225
- Name Server zones 252
- Name Server Zoning 238
- Name Server zoning rules 242
- Name Service Table 230
- naming convention 159
- Netscape 77

- nicknames 122
- non-RAID 109
- notification options 85
- Novell NetWare FC attachment 26
- NUMA-Q 15

O

- Open System Storage 104
- orphan zone 245, 256
- OS/400 operating system 11

P

- parallel interface 6
- path group 5
- Point to Point 308
- port granularity 243
- port groups 238
- PPRC 77
- Private Attributes 364
- Private Initiator 232
- Private loop 229
- private loop 229
- private loop initiators 232
- Private Target 232
- private target 233
- Problem Log 84
- Problem Notification 85
- pSeries attachment requirements 12
- pSeries FC host 11
- pseudo-host 163
- Public loop 229
- public loop 229

Q

- QLA2100F 26
- QLA2200F 26–27
- QLogic QLA2100F 26, 47–48, 58–59
- QLogic QLA2200F 27, 38, 43, 47, 49, 58, 60
- queue depth 24
- Quick Initialize 345

R

- RAID 109
- RAID 5 75, 109, 115, 130
- Redbooks Web site 368
 - Contact us xii
- Remove logical drive 357

- Re-Scan SCSI Bus 295
- Router config file 328
- Router LED codes 323
- Router Node Mapping 323
- Router power up sequence 325
- Router Properties 334
- Router Subsystem Diagnostic test 323
- RS/6000 11

S

- SAN Data Gateway 24, 271
- SCSI 271, 276, 283, 294, 318
- SCSI Channel 286
- SCSI device map 295
- SCSI to LUN map 304
- scsiRescan 295
- SDD 143, 318
- security 237
- segmentation 238
- Server 277
- Service port 273
- setting a Gateway address 275
- setting a Subnetmask 275
- Setting the Ethernet Address 275
- setting the IP address 275
- short-wave ports 273
- SignOn Drive 336
- Simple Network Management Protocol 85
- SLIC Manager 363
- SLIC Manager daemon 326, 331, 363
- SLIC Manager software 326
- SLIC Zone 327
- SLIC zone 336, 363
- SMIT 14
- SNMP 85, 266
- sort criteria 144
- spreading volumes 132
- SSA 321, 334, 361
- SSA loops 115
- Starting the SLIC Manager 332
- Startup sequence 276
- Storage Allocation 89, 96
- Storage Server Attributes 162
- storage type 130
- StorWatch Enterprise Storage Server Expert 119
- StorWatch SAN Data Gateway Specialist 277, 280
- Subsystem Device Driver 44, 143, 317
- Sun 24, 45

Sun FC attachment 30
switch cascading 8
Switch registration 309

T

T_Port 264
T_Ports 264
Target Hosts 144
Technology Pack 17
Telnet 275
TFTP 266
TimeOutValue 53, 64
Tivoli SANergy 364
TL_Port 235, 243
TL_Ports 235
TL_Ports zoning 236
TL-CFG 244
TL-Cfg 235
topology 167
Track Format 109
track format 108
Translation Entries List 237
translation entries list 236
translative device addresses 236
trivial file transfer protocol 266
trunk ports 264

U

Uniform Resource Locator 77
unmap 359
UnMapped 360
URL 77
Users 93

V

Vicom Fibre Channel SLIC Router 321
Virtual Private SAN 297, 303, 306, 315
Virtual Private SAN (VPS) 296
volume assignment 127
volume definitions 131
volume placement 128
volumes sequentially placed 132
VPS 299, 309
VPS Registration Service 298

W

Windows 2000 FC attachment 58

Windows NT 4.0 FC attachment 47
worldwide port name 16
WWN 229, 234, 242, 263
WWNN 79, 166
WWPN 4, 97, 118, 122, 159, 163
WWPN naming convention 164
WWPN value 69

X

xSeries FC attachment 15

Z

Zoning 296, 306, 316
zoning 273, 319



Implementing an Open IBM SAN Featuring the INRANGE Portfolio

(0.5" spine)
0.475" <-> 0.875"
250 <-> 459 pages



Implementing an Open IBM SAN

Featuring the INRANGE Portfolio

Discover the latest additions to the IBM SAN family

Enhance your skills while using an easy-to-follow format

Grow with the new technology

"Do everything that is necessary and absolutely nothing that is not."

In this IBM Redbook, we have tried to consolidate as much of the critical information as possible while covering procedures and tasks that are likely to be encountered on a daily basis.

Each of the products described has much, much more functionality than we could ever hope to cover in just one redbook. The IBM SAN portfolio is rich in quality products that bring a vast amount of technicality and vitality to the SAN world. Their inclusion and selection is based on a thorough understanding of the storage networking environment that positions IBM, and therefore its customers and partners, in an ideal position to take advantage by their deployment.

In this redbook we cover the latest additions to the IBM SAN family, focusing on the products from INRANGE. We show how they can be implemented in an open systems environment, and we focus on the Fibre Channel protocol (FCP) environment in particular. We address some of the key concepts that they bring to the market, and in each case, we give an overview of those functions that are essential to building a robust SAN environment.

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

For more information:
ibm.com/redbooks